



Final Report

National Security Commission on Artificial Intelligence





NATIONAL
SECURITY
COMMISSION
ON ARTIFICIAL
INTELLIGENCE

COMMISSION MEMBERS

Eric Schmidt
Chair

Robert Work
Vice Chair

Safra Catz

Eric Horvitz

Steve Chien

Andrew Jassy

Mignon Clyburn

Gilman Louie

Chris Darby

William Mark

Kenneth Ford

Jason Matheny

José-Marie Griffiths

Katharina McFarland

Andrew Moore

Letter from the Chair and Vice Chair

Americans have not yet grappled with just how profoundly the artificial intelligence (AI) revolution will impact our economy, national security, and welfare. Much remains to be learned about the power and limits of AI technologies. Nevertheless, big decisions need to be made now to accelerate AI innovation to benefit the United States and to defend against the malign uses of AI.

When considering these decisions, our leaders confront the classic dilemma of statecraft identified by Henry Kissinger: “When your scope for action is greatest, the knowledge on which you can base this action is always at a minimum. When your knowledge is greatest, the scope for action has often disappeared.” The scope for action remains, but America’s room for maneuver is shrinking.

As a bipartisan commission of 15 technologists, national security professionals, business executives, and academic leaders, the National Security Commission on Artificial Intelligence (NSCAI) is delivering an uncomfortable message: America is not prepared to defend or compete in the AI era. This is the tough reality we must face. And it is this reality that demands comprehensive, whole-of-nation action. Our final report presents a strategy to defend against AI threats, responsibly employ AI for national security, and win the broader technology competition for the sake of our prosperity, security, and welfare. The U.S. government cannot do this alone. It needs committed partners in industry, academia, and civil society. And America needs to enlist its oldest allies and new partners to build a safer and freer world for the AI era.

AI is an inspiring technology. It will be the most powerful tool in generations for benefiting humanity. Scientists have already made astonishing progress in fields ranging from biology and medicine to astrophysics by leveraging AI. These advances are not science fair experiments; they are improving life and unlocking mysteries of the natural world. They are the kind of discoveries for which the label “game changing” is not a cliché.

AI systems will also be used in the pursuit of power. We fear AI tools will be weapons of first resort in future conflicts. AI will not stay in the domain of superpowers or the realm of science fiction. AI is dual-use, often open-source, and diffusing rapidly. State adversaries

are already using AI-enabled disinformation attacks to sow division in democracies and jar our sense of reality. States, criminals, and terrorists will conduct AI-powered cyber attacks and pair AI software with commercially available drones to create “smart weapons.” It is no secret that America’s military rivals are integrating AI concepts and platforms to challenge the United States’ decades-long technology advantage. We will not be able to defend against AI-enabled threats without ubiquitous AI capabilities and new warfighting paradigms. We want the men and women in national security departments and agencies to have access to the best technology in the world to defend themselves and us, and to protect our interests and those of our allies and partners.

Despite exciting experimentation and a few small AI programs, the U.S. government is a long way from being “AI-ready.” The Commission’s business leaders are most frustrated by slow government progress because they know it’s possible for large institutions to adopt AI. AI integration is hard in any sector—and the national security arena poses some unique challenges. Nevertheless, committed leaders can drive change. We need those leaders in the Pentagon and across the Federal Government to build the technical infrastructure and connect ideas and experimentation to new concepts and operations. By 2025, the Department of Defense and the Intelligence Community must be AI-ready.

We should embrace the AI competition. Competition already infuses the quests for data, computing power, and the holy grail: the rare talent to make AI breakthroughs. The fact that AI courses through so many adjacent technologies and is leveraged across so many fields explains its power and leads inexorably to another critical point: AI is part of a broader global technology competition. Competition will speed up innovation. We should race together with partners when AI competition is directed at the moonshots that benefit humanity like discovering vaccines. But we must win the AI competition that is intensifying strategic competition with China. China’s plans, resources, and progress should concern all Americans. It is an AI peer in many areas and an AI leader in some applications. We take seriously China’s ambition to surpass the United States as the world’s AI leader within a decade.

The AI competition is also a values competition. China’s domestic use of AI is a chilling precedent for anyone around the world who cherishes individual liberty. Its employment of AI as a tool of repression and surveillance—at home and, increasingly, abroad—is a powerful counterpoint to how we believe AI should be used. The AI future can be democratic, but we have learned enough about the power of technology to strengthen authoritarianism abroad and fuel extremism at home to know that we must not take for granted that future technology trends will reinforce rather than erode democracy. We must work with fellow democracies and the private sector to build privacy-protecting standards into AI technologies and advance democratic norms to guide AI uses so that democracies can responsibly use AI tools for national security purposes.

We would like to emphasize a few areas where action is necessary because the stakes of the competition are so high:

Leadership.

Ultimately, we have a duty to convince the leaders in the U.S. Government to make the hard decision and the down payment to win the AI era. In America, the buck stops with the President, and AI strategy starts in the White House. We built a National Security Council to confront the challenges of the post–World War II era. Now we need to create a Technology Competitiveness Council to build a strategy that accounts for the complex security, economic, and scientific challenges of AI and its associated technologies. That leadership imperative extends into all critical national security departments and agencies.

Talent.

The human talent deficit is the government's most conspicuous AI deficit and the single greatest inhibitor to buying, building, and fielding AI-enabled technologies for national security purposes. This is not a time to add a few new positions in national security departments and agencies for Silicon Valley technologists and call it a day. We need to build entirely new talent pipelines from scratch. We should establish a new Digital Service Academy and civilian National Reserve to grow tech talent with the same seriousness of purpose that we grow military officers. The digital age demands a digital corps. Just as important, the United States needs to win the international talent competition by improving both STEM education and our system for admitting and retaining highly skilled immigrants.

Hardware.

Microelectronics power all AI, and the United States no longer manufactures the world's most sophisticated chips. We do not want to overstate the precariousness of our position, but given that the vast majority of cutting-edge chips are produced at a single plant separated by just 110 miles of water from our principal strategic competitor, we must reevaluate the meaning of supply chain resilience and security. A recent chip shortage for auto manufacturing cost an American car company an estimated \$2.5 billion. A strategic blockage would cost far more and put our security at risk. The federal investment and incentives needed to revitalize domestic microchip fabrication—perhaps \$35 billion—should be an easy decision when the alternative is relying on another country to produce the engines that power the machines that will shape the future.

Innovation Investment.

We worry that only a few big companies and powerful states will have the resources to make the biggest AI breakthroughs. Despite the diffusion of open-source tools, the needs for computing power and troves of data to improve algorithms are soaring at the cutting edge of innovation. The federal government must partner with U.S. companies to preserve American leadership and to support development of diverse AI applications that advance the national interest in the broadest sense. If anything, this report underplays the investments America will need to make. The \$40 billion we recommend to expand and democratize federal AI research and development (R&D) is a modest down payment on future breakthroughs. We will also need to build secure digital infrastructure across the nation, shared cloud computing access, and smart cities to truly leverage AI for the benefit of all Americans. We envision hundreds of billions in federal spending in the coming years.

This is not a time for abstract criticism of industrial policy or fears of deficit spending to stand in the way of progress. In 1956, President Dwight Eisenhower, a fiscally conservative Republican, worked with a Democratic Congress to commit \$10 billion to build the Interstate Highway System. That is \$96 billion in today's world. Surely we can make a similar investment in the nation's future.

We are proud of the NSCAI's bipartisan work. We have debated together, learned together, and achieved consensus on critical points. It is our privilege to submit our recommendations to Congress and the President. To paraphrase Winston Churchill, we are at the beginning of the beginning of the competition that will shape our prosperity, national security, and the well-being of our citizens. Our report presents the first steps the United States should take to defend, compete, and win in the AI era.



Eric Schmidt,
Chair



Bob Work,
Vice Chair

LETTER FROM THE EXECUTIVE DIRECTOR:

The Beginning of the Beginning

When we started our journey two years ago, little did we know what was in front of us. What we encountered was willingness and hope among many friends and allies to get our mission from Congress right to maintain the United States' advantage in artificial intelligence (AI).

We enjoyed support from U.S. Departments and Agencies. Many of them loaned us resources, including detailing both civilian and military personnel, and dedicated countless hours to help us understand their missions and priorities. Members of Congress and congressional staff worked closely with us to accelerate our government's adoption of AI for national security purposes.

Over the course of the Commission's work, we engaged with hundreds of representatives from the private sector, academia, civil society, and across the government. We received countless briefings—classified and unclassified. We met with anyone who thinks about AI, works with AI, and develops AI who was willing to make time for us.

We found consensus among nearly all of our partners on three points: the conviction that AI is an enormously powerful technology, acknowledgement of the urgency to invest more in AI innovation, and responsibility to develop and use AI guided by democratic principles.

We also talked to our allies—old and new. From New Delhi to Tel Aviv to London, there was a willingness and desire to work with the United States to deepen cooperation on AI.

I am indebted to the many individuals who volunteered with us, interned with us, provided expertise, and were friends of the Commission. I am particularly grateful to the dedicated full-time staff of the Commission, who in many cases stepped away from important jobs to join this essential mission.

In the last two years, we encountered widespread hope that AI could generate incredible benefits for our nation's economy, welfare, and security. We also heard concern that AI—like any technology—could create new challenges and exacerbate existing problems. We listened and took those concerns seriously.

We ultimately came away with a recognition that if America embraces and invests in AI based on our values, it will transform our country and ensure that the United States and its allies continue to shape the world for the good of all humankind.

Thank you!
Yll Bajraktari

Executive Summary

No comfortable historical reference captures the impact of artificial intelligence (AI) on national security. AI is not a single technology breakthrough, like a bat-wing stealth bomber. The race for AI supremacy is not like the space race to the moon. AI is not even comparable to a general-purpose technology like electricity. However, what Thomas Edison said of electricity encapsulates the AI future: “It is a field of fields . . . it holds the secrets which will reorganize the life of the world.” Edison’s astounding assessment came from humility. All that he discovered was “very little in comparison with the possibilities that appear.”

The National Security Commission on Artificial Intelligence (NSCAI) humbly acknowledges how much remains to be discovered about AI and its future applications. Nevertheless, we know enough about AI today to begin with two convictions.

First, the rapidly improving ability of computer systems to solve problems and to perform tasks that would otherwise require human intelligence—and in some instances exceed human performance—is world altering. AI technologies are the most powerful tools in generations for expanding knowledge, increasing prosperity, and enriching the human experience. AI is also the quintessential “dual-use” technology. The ability of a machine to perceive, evaluate, and act more quickly and accurately than a human represents a competitive advantage in any field—civilian or military. AI technologies will be a source of enormous power for the companies and countries that harness them.

Second, AI is expanding the window of vulnerability the United States has already entered. For the first time since World War II, America’s technological predominance—the backbone of its economic and military power—is under threat. China possesses the might, talent, and ambition to surpass the United States as the world’s leader in AI in the next decade if current trends do not change. Simultaneously, AI is deepening the threat posed by cyber attacks and disinformation campaigns that Russia, China, and others are using to infiltrate our society, steal our data, and interfere in our democracy. The limited uses of AI-enabled attacks to date represent the tip of the iceberg. Meanwhile, global crises exemplified by the COVID-19 pandemic and climate change highlight the need to expand our conception of national security and find innovative AI-enabled solutions.



“The NSCAI Final Report presents an integrated national strategy to reorganize the government, reorient the nation, and rally our closest allies and partners to defend and compete in the coming era of AI-accelerated competition and conflict.”

Given these convictions, the Commission concludes that the United States must act now to field AI systems and invest substantially more resources in AI innovation to protect its security, promote its prosperity, and safeguard the future of democracy. Today, the government is not organizing or investing to win the technology competition against a committed competitor, nor is it prepared to defend against AI-enabled threats and rapidly adopt AI applications for national security purposes. This is not a time for incremental toggles to federal research budgets or adding a few new positions in the Pentagon for Silicon Valley technologists. This will be expensive and require a significant change in mindset. America needs White House leadership, Cabinet-member action, and bipartisan Congressional support to win the AI era.

The NSCAI Final Report presents an integrated national strategy to reorganize the government, reorient the nation, and rally our closest allies and partners to defend and compete in the coming era of AI-accelerated competition and conflict. It is a two-pronged approach. Part I, “Defending America in the AI Era,” outlines the stakes, explains what the United States must do to defend against the spectrum of AI-related threats, and recommends how the U.S. government can responsibly use AI technologies to protect the American people and our interests. Part II, “Winning the Technology Competition,” addresses the critical elements of the AI competition and recommends actions the government must take to promote AI innovation to improve national competitiveness and protect critical U.S. advantages. The recommendations are designed as interlocking and mutually reinforcing actions that must be taken together.

Part I: Defending America in the AI Era.

AI-enhanced capabilities will be the tools of first resort in a new era of conflict as strategic competitors develop AI concepts and technologies for military and other malign uses and cheap and commercially available AI applications ranging from “deepfakes” to lethal drones become available to rogue states, terrorists, and criminals. The United States must prepare to defend against these threats by quickly and responsibly adopting AI for national security and defense purposes. Defending against AI-capable adversaries operating at machine speeds without employing AI is an invitation to disaster. Human operators will not be able to keep up with or defend against AI-enabled cyber or disinformation attacks, drone swarms, or missile attacks without the assistance of AI-enabled machines. National security professionals must have access to the world’s best technology to protect themselves, perform their missions, and defend us. The Commission recommends that the government take the following actions:

Defend against emerging AI-enabled threats to America’s free and open society. Digital dependence in all walks of life is transforming personal and commercial vulnerabilities into potential national security weaknesses. Adversaries are using AI systems to enhance disinformation campaigns and cyber attacks. They are harvesting data on Americans to build profiles of their beliefs, behavior, and biological makeup for tailored attempts to manipulate or coerce individuals. This gathering storm of foreign influence and interference requires organizational and policy reforms to bolster our resilience. The government needs to stand up a task force and 24/7 operations center to confront digital disinformation. It needs to better secure its own databases and prioritize data security in foreign investment screening, supply chain risk management, and national data protection legislation. The government should leverage AI-enabled cyber defenses to protect against AI-enabled cyber attacks. And biosecurity must become a top-tier priority in national security policy.

Prepare for future warfare. Our armed forces’ competitive military-technical advantage could be lost within the next decade if they do not accelerate the adoption of AI across their missions. This will require marrying top-down leadership with bottom-up innovation to put operationally relevant AI applications into place. The Department of Defense (DoD) should:

First, establish the foundations for widespread integration of AI by 2025. This includes building a common digital infrastructure, developing a digitally-literate workforce, and instituting more agile acquisition, budget, and oversight processes. It also requires strategically divesting from military systems that are ill-equipped for AI-enabled warfare and instead investing in next-generation capabilities.

Second, achieve a state of military AI readiness by 2025. Pentagon leadership must act now to drive organizational reforms, design innovative warfighting concepts, establish AI and digital readiness performance goals, and define a joint warfighting network

architecture. DoD must also augment and focus its AI R&D portfolio. Readiness will also require promoting AI interoperability with allies and partners.

Manage risks associated with AI-enabled and autonomous weapons. AI will enable new levels of performance and autonomy for weapon systems. But it also raises important legal, ethical, and strategic questions surrounding the use of lethal force. Provided their use is authorized by a human commander or operator, properly designed and tested AI-enabled and autonomous weapon systems can be used in ways that are consistent with international humanitarian law. DoD's rigorous, existing weapons review and targeting procedures, including its dedicated protocols for autonomous weapon systems and commitment to strong AI ethical principles, are capable of ensuring that the United States will field safe and reliable AI-enabled and autonomous weapon systems and use them in a lawful manner. While it is neither feasible nor currently in the interests of the United States to pursue a global prohibition of AI-enabled and autonomous weapon systems, the global, unchecked use of such systems could increase risks of unintended conflict escalation and crisis instability. To reduce the risks, the United States should (1) clearly and publicly affirm existing U.S. policy that only human beings can authorize employment of nuclear weapons and seek similar commitments from Russia and China; (2) establish venues to discuss AI's impact on crisis stability with competitors; and (3) develop international standards of practice for the development, testing, and use of AI-enabled and autonomous weapon systems.

Transform national intelligence. The Intelligence Community (IC) should adopt and integrate AI-enabled capabilities across all aspects of its work, from collection to analysis. Intelligence will benefit from AI more than any other national security mission. To capitalize on AI, the Office of the Director of National Intelligence needs to empower and resource its science and technology leaders. The entire IC should leverage open-source and publicly available information in its analysis and prioritize collection of scientific and technical intelligence. For better insights, intelligence agencies will need to develop innovative approaches to human-machine teaming that use AI to augment human judgment.

Scale up digital talent in government. National security agencies need more digital experts now or they will remain unprepared to buy, build, and use AI and associated technologies. The talent deficit in DoD and the IC represents the greatest impediment to being AI-ready by 2025. The government needs new talent pipelines, including a U.S. Digital Service Academy to train current and future employees. It needs a civilian National Digital Reserve Corps to recruit people with the right skills—including industry experts, academics, and recent college graduates. And it needs a Digital Corps, modeled on the Army Medical Corps, to organize technologists already serving in government.

Establish justified confidence in AI systems. If AI systems routinely do not work as designed or are unpredictable in ways that can have significant negative consequences, then leaders will not adopt them, operators will not use them, Congress will not fund them, and the

American people will not support them. To establish justified confidence, the government should focus on ensuring that its AI systems are robust and reliable, including through research and development (R&D) investments in AI security and advancing human-AI teaming through a sustained initiative led by the national research labs. It should also enhance DoD's testing and evaluation capabilities as AI-enabled systems grow in number, scope, and complexity. Senior-level responsible AI leads should be appointed across the government to improve executive leadership and policy oversight.

Present a democratic model of AI use for national security. AI tools are critical for U.S. intelligence, homeland security, and law enforcement agencies. Public trust will hinge on justified assurance that government use of AI will respect privacy, civil liberties, and civil rights. The government must earn that trust and ensure that its use of AI tools is effective, legitimate, and lawful. This imperative calls for developing AI tools to enhance oversight and auditing, increasing public transparency about AI use, and building AI systems that advance the goals of privacy preservation and fairness. It also requires ensuring that those impacted by government actions involving AI can seek redress and have due process. The government should strengthen oversight and governance mechanisms and establish a task force to assess evolving concerns about AI and privacy, civil liberties, and civil rights.

Part II: Winning the Technology Competition.

The race to research, develop, and deploy AI and associated technologies is intensifying the technology competition that underpins a wider strategic competition. China is organized, resourced, and determined to win this contest. The United States retains advantages in critical areas, but current trends are concerning. While a competitive response is complicated by deep academic and commercial interconnections, the United States must do what it takes to retain its innovation leadership and position in the world. The U.S. government must embrace the AI competition and organize to win it by orchestrating and aligning U.S. strengths.

Organize with a White House–led strategy for technology competition. The United States must elevate AI considerations from the technical to the strategic level. Emerging technologies led by AI now underpin our economic prosperity, security, and welfare. The White House should establish a new Technology Competitiveness Council led by the Vice President to integrate security, economic, and scientific considerations; develop a comprehensive technology strategy; and oversee its implementation.

Win the global talent competition. The United States risks losing the global competition for scarce AI expertise if it does not cultivate more potential talent at home and recruit and retain more existing talent from abroad. The United States must move aggressively on both fronts. Congress should pass a National Defense Education Act II to address deficiencies across the American educational system—from K-12 and job reskilling to investing in

thousands of undergraduate- and graduate-level fellowships in fields critical to the AI future. At the same time, Congress should pursue a comprehensive immigration strategy for highly skilled immigrants to encourage more AI talent to study, work, and remain in the United States through new incentives and visa, green card, and job-portability reforms.

Accelerate AI innovation at home. The government must make major new investments in AI R&D and establish a national AI research infrastructure that democratizes access to the resources that fuel AI development across the nation. The government should: (1) double non-defense funding for AI R&D annually to reach \$32 billion per year by 2026, establish a National Technology Foundation, and triple the number of National AI Research Institutes; (2) establish a National AI Research Infrastructure composed of cloud computing resources, test beds, large-scale open training data, and an open knowledge network that will broaden access to AI and support experimentation in new fields of science and engineering; and (3) strengthen commercial competitiveness by creating markets for AI and by forming a network of regional innovation clusters.

Implement comprehensive intellectual property (IP) policies and regimes. The United States must recognize IP policy as a national security priority critical for preserving America's leadership in AI and emerging technologies. This is especially important in light of China's efforts to leverage and exploit IP policies. The United States lacks the comprehensive IP policies it needs for the AI era and is hindered by legal uncertainties in current U.S. patent eligibility and patentability doctrine. The U.S. government needs a plan to reform IP policies and regimes in ways that are designed to further national security priorities.

Build a resilient domestic base for designing and fabricating microelectronics. After decades leading the microelectronics industry, the United States is now almost entirely reliant on foreign sources for production of the cutting-edge semiconductors that power all the AI algorithms critical for defense systems and everything else. Put simply: the U.S. supply chain for advanced chips is at risk without concerted government action. Rebuilding domestic chip manufacturing will be expensive, but the time to act is now. The United States should commit to a strategy to stay at least two generations ahead of China in state-of-the-art microelectronics and commit the funding and incentives to maintain multiple sources of cutting-edge microelectronics fabrication in the United States.

Protect America's technology advantages. As the margin of U.S. technological advantage narrows and foreign efforts to acquire American know-how and dual-use technologies increase, the United States must reexamine how to best protect ideas, technology, and companies without unduly hindering innovation. The United States must:

First, modernize export controls and foreign investment screening to better protect critical dual-use technologies—including by building regulatory capacity and fully implementing recent legislative reforms, implementing coordinated export controls on advanced semiconductor manufacturing equipment with allies, and expanding disclosure requirements for investors from competitor nations.

Second, protect the U.S. research enterprise as a national asset—by providing government agencies, law enforcement, and research institutions with tools and resources to conduct nuanced risk assessments and share information on specific threats and tactics, coordinating research protection efforts with allies and partners, bolstering cybersecurity support for research institutions, and strengthening visa vetting to limit problematic research collaborations.

Build a favorable international technology order. The United States must work hand-in-hand with allies and partners to promote the use of emerging technologies to strengthen democratic norms and values, coordinate policies and investments to advance global adoption of digital infrastructure and technologies, defend the integrity of international technical standards, cooperate to advance AI innovation, and share practices and resources to defend against malign uses of technology and the influence of authoritarian states in democratic societies. The United States should lead an Emerging Technology Coalition to achieve these goals and establish a Multilateral AI Research Institute to enhance the United States' position as a global research hub for emerging technology. The Department of State should be reoriented, reorganized, and resourced to lead diplomacy in emerging technologies.

Win the associated technologies competitions. Leadership in AI is necessary but not sufficient for overall U.S. technological leadership. AI sits at the center of the constellation of emerging technologies, enabling some and enabled by others. The United States must therefore develop a single, authoritative list of the technologies that will underpin national competitiveness in the 21st century and take bold action to catalyze U.S. leadership in AI, microelectronics, biotechnology, quantum computing, 5G, robotics and autonomous systems, additive manufacturing, and energy storage technology. U.S. leadership across these technologies requires investing in specific platforms that will enable transformational breakthroughs and building vibrant domestic manufacturing ecosystems in each. At the same time, the government will need to continuously identify and prioritize emerging technologies farther over the horizon.

Conclusion

This new era of competition promises to change the world we live in and how we live within it. We can either shape the change to come or be swept along by it. We now know that the uses of AI in all aspects of life will grow and the pace of innovation will continue to accelerate. We know adversaries are determined to turn AI capabilities against us. We know China is determined to surpass us in AI leadership. We know advances in AI build on themselves and confer significant first-mover advantages. Now we must act. The principles we establish, the federal investments we make, the national security applications we field, the organizations we redesign, the partnerships we forge, the coalitions we build, and the talent we cultivate will set America's strategic course. The United States should invest what it takes to maintain its innovation leadership, to responsibly use AI to defend free people and free societies, and to advance the frontiers of science for the benefit of all humanity. AI is going to reorganize the world. America must lead the charge.

Preface

The National Security Commission on Artificial Intelligence's (NSCAI) task is to make recommendations to the President and Congress to “advance the development of artificial intelligence [AI], machine learning, and associated technologies to comprehensively address the national security and defense needs of the United States.” In establishing the Commission, Section 1051 of the John S. McCain National Defense Authorization Act for Fiscal Year 2019 instructs NSCAI to examine AI through the lenses of national competitiveness, the means to sustain technological advantage, trends in international cooperation and competitiveness, ways to foster investment in basic and advanced research, workforce and training, potential risks of military use, ethical concerns, establishment of data standards and incentivization of data sharing, and the future evolution of AI.¹

The 15 commissioners were nominated by Congress and the Executive Branch. They represent a diverse group of technologists, business executives, academic leaders, and national security professionals. They have approached all inquiries in bipartisan fashion and reached consensus on the Final Report. The Commission's operations have been guided by two principles: the need for action and the importance of transparency.

Action.

The Commission's work includes an initial report in July 2019, interim reports in November 2019 and October 2020, two additional quarterly memorandums, a series of special papers in response to the COVID-19 pandemic, and now a final report. Waiting to deliver recommendations in a final report was not an option when we began our work in the spring of 2019. Assessing the broad national security implications of a dynamic technology like AI at a single point in time is like trying to catch lightning in a bottle. Scientists continue to deliver AI breakthroughs and the commercial sector is finding new ways to apply AI at an accelerating pace. Competitors around the world are developing AI strategies and investing resources. The Commission delivered recommendations on a continuous basis, aiming to match the speed of AI developments and the desires from the Executive Branch and Congress for help in deciding what to do. Congress has already adopted a number of our recommendations in the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021,² and the Executive Branch has incorporated recommendations

as well. And we have continuously sought to learn from and educate a wide range of stakeholders to build a shared understanding about how AI will impact national security.

Transparency.

The NSCAI has been committed to transparency. As a Federal Advisory Committee, it has held five public plenary sessions totaling approximately 15 hours of deliberations, streamed live online, and archived meeting recordings on the NSCAI website. It has responded to more than two dozen Freedom of Information Act requests and released more than 2,500 pages of material. NSCAI has posted more than 700 pages of draft materials for public review and comment. With the exception of materials and issues classified for national security reasons, the Commission has endeavored to offer full transparency. We have proactively engaged with the media after every plenary session, quarterly report, and submission to Congress. In dozens of separate engagements, we have partnered with non-governmental organizations, federal government organizations, and international organizations to communicate our recommendations to the media and the public.

Most important, we have taken on the hardest issues with AI in public settings and made recommendations only after consulting with a wide range of civil society, private sector, and government groups. We have tried to listen and understand views across the spectrum on deeply complicated aspects of AI. We have engaged ethicists, technologists, and national security strategists. We have spoken with warriors and diplomats. We have talked to academics and entrepreneurs. All told, commissioners and staff have participated in hundreds of discussions. As the commissioners built consensus on recommendations, we approached issues with care and humility.

The Final Report.

The Final Report presents the NSCAI's recommendations as a strategy for winning the AI era. The 16 chapters in the Main Report provide topline recommendations. The accompanying Blueprints for Action outline concrete steps that departments and agencies can take to implement NSCAI recommendations. The Commission has provided as much specificity as possible—including by providing draft legislative text and executive orders—to help the President and Congress move rapidly from understanding AI to acting for the benefit of the American people.

The Final Report represents an important step, but it is not the NSCAI's final act. For the remaining life of the Commission, our work will focus on implementation to help the President and Congress make the investments and take the actions recommended to win the AI era.

¹ For full text, see Pub. L. 115-232, 132 Stat. 1636 (2018), <https://www.congress.gov/115/bills/hr5515/BILLS-115hr5515enr.pdf>.

² For full text, see Pub. L. 116-283, 134 Stat. 3388 (2021), <https://www.congress.gov/bill/116th-congress/house-bill/6395/text>.

Table of Contents

Introduction	19
Artificial Intelligence in Context	31
PART I: DEFENDING AMERICA IN THE AI ERA	41
Chapter 1: Emerging Threats in the AI Era	43
Chapter 2: Foundations of Future Defense	59
Chapter 3: AI and Warfare	75
Chapter 4: Autonomous Weapon Systems and Risks Associated with AI-Enabled Warfare	89
Chapter 5: AI and the Future of National Intelligence	107
Chapter 6: Technical Talent in Government	119
Chapter 7: Establishing Justified Confidence in AI Systems	131
Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security	141
PART II: WINNING THE TECHNOLOGY COMPETITION	155
Chapter 9: A Strategy for Competition and Cooperation	157
Chapter 10: The Talent Competition	171
Chapter 11: Accelerating AI Innovation	183
Chapter 12: Intellectual Property	199
Chapter 13: Microelectronics	211
Chapter 14: Technology Protection	223
Chapter 15: A Favorable International Technology Order	241
Chapter 16: Associated Technologies	253
Blueprints for Action	271
Appendices	599

Introduction

Artificial Intelligence (AI) technologies promise to be the most powerful tools in generations for expanding knowledge, increasing prosperity, and enriching the human experience. The technologies will be the foundation of the innovation economy and a source of enormous power for countries that harness them. AI will fuel competition between governments and companies racing to field it. And it will be employed by nation-states to pursue their strategic ambitions.

Americans have not yet seriously grappled with how profoundly the AI revolution will impact society, the economy, and national security. Recent AI breakthroughs, such as a computer defeating a human in the popular strategy game of Go¹, shocked other nations into action, but it did not inspire the same response in the United States. Despite our private-sector and university leadership in AI, the United States remains unprepared for the coming era. Americans must recognize the assertive role that the government will have to play in ensuring the United States wins this innovation competition. Congress and the President will have to support the scale of public resources required to achieve it.

The magnitude of the technological opportunity coincides with a moment of strategic vulnerability. China is a competitor possessing the might, talent, and ambition to challenge America's technological leadership, military superiority, and its broader position in the world. AI is deepening the threat posed by cyber attacks and disinformation campaigns that Russia, China, and other state and non-state actors are using to infiltrate our society, steal our data, and interfere in our democracy. The limited uses of AI-enabled attacks to date are the tip of the iceberg. Meanwhile, global crises exemplified in the global pandemic and climate change are expanding the definition of national security and crying out for innovative technological solutions. AI can help us navigate many of these new challenges.

We are fortunate. The AI revolution is not a strategic surprise. We are experiencing its impact in our daily lives and can anticipate how research progress will translate into real-world applications before we have to confront the full national security ramifications. This commission can warn of national security challenges and articulate the benefits, rather than explain why previous warnings were ignored and opportunities were missed. We still have a window to make the changes to build a safer and better future. The pace of AI innovation is not flat; it is accelerating. If the United States does not act, it will likely lose its leadership position in AI to China in the next decade and become more vulnerable to a spectrum of AI-enabled threats from a host of state and non-state actors.

The Commission concludes that the United States needs to implement a strategy to defend and compete in the AI era. The White House must lead the effort to reorganize the government and reorient the nation. This report presents the core elements of the strategy.

- Part I, “Defending America in the AI Era” (Chapters 1–8), outlines what the United States must do to defend against the spectrum of AI-related threats from state and non-state actors and recommends how the U.S. government can responsibly use AI technologies to protect the American people and our interests.
- Part II, “Winning the Technology Competition” (Chapters 9–16), outlines AI’s role in a broader technology competition. Each chapter addresses a critical element of the competition and recommends actions the government must take to promote AI innovation to improve national competitiveness and protect critical U.S. advantages.

Why Does AI Matter?

In 1901, Thomas Edison was asked to predict electricity’s impact on humanity. Two decades after the development of the light bulb, he foresaw a general-purpose technology of unlimited possibilities. “[Electricity] is the field of fields,” he said. “It holds the secrets which will reorganize the life of the world.”² AI is a very different kind of general-purpose technology, but we are standing at a similar juncture and see a similarly wide-ranging impact.³ The rapidly improving ability of computer systems to solve problems and to perform tasks that would otherwise require human intelligence is transforming many aspects of human life and every field of science. It will be incorporated into virtually all future technology. The entire innovation base supporting our economy and security will leverage AI. How this “field of fields” is used—for good and for ill—will reorganize the world.

The Commission’s assessment is rooted in a realistic understanding of AI’s current state of development and a projection of how the technology will evolve.

AI is already ubiquitous in everyday life and the pace of innovation is accelerating. We take for granted that AI already shapes our lives in ways small and big. A “smartphone” has multiple AI-enabled features including voice assistants, photo tagging, facial recognition security, search apps, recommendation and advertising engines, and less obvious AI enhancements in its operating system. AI is helping predict the spread and escalation of a pandemic outbreak, planning and optimizing the distribution of goods and services, monitoring traffic flow and safety, speeding up drug and therapeutic discovery, and automating routine office functions. Recognizing the pace of change is critical to understanding the power of AI. The application of AI techniques to solve problems is compressing innovation timescales and turning once-fantastical ideas into realities across a range of disciplines.

Deploying and adopting AI remains a hard problem. AI cannot magically solve problems. As AI moves from an elite niche science to a mainstream tool, engineering will be as important as scientific breakthroughs. Early adopters across sectors have learned similar lessons: Trying to employ AI is a slog even after the science is settled. Many of the most important real-world impacts will come from figuring out how to employ existing AI algorithms and systems, some more than a decade old. The integration challenge is immense. Harnessing data, hardening and packaging laboratory algorithms so they are ready for use in the field, and adapting AI software to legacy equipment and rigid organizations all require time, effort, and patience. Integrating AI often necessitates overcoming substantial organizational and cultural barriers, and it demands top-down leadership.

AI tools are diffusing broadly and rapidly. Cutting-edge deep learning techniques are often prohibitively expensive, requiring vast amounts of data, computing power, and specialized knowledge. However, AI will not be the provenance of only big states and big tech. Many machine learning tools that fuel AI applications are publicly available and usable even for non-experts. Open-source applications and development tools combined with inexpensive cloud computing and less data-intensive approaches are expanding AI opportunities across the world to state and non-state actors.

AI is changing relationships between humans and machines. In modern society, we already rely much more on machines and automation than we may be aware. The U.S. military, for instance, has used autonomous systems for decades. However, as AI capabilities improve, the dynamics within human-machine “teams” will change. In the past, computers could only perform tasks that fell within a clearly defined set of parameters or rules programmed by a human. As AI becomes more capable, computers will be able to learn and perform tasks based on parameters that humans do not explicitly program, creating choices and taking actions at a volume and speed never before possible. Across many fields of human activity, AI innovations are raising important questions about what choices to delegate to intelligent machines, in what circumstances, and for what reasons. In the national security sphere, these questions will take on greater significance as AI is integrated into defense and intelligence systems. Across our entire society, we will need to address these new complexities with nuanced approaches, intellectual curiosity, and care that recognizes the increasing ubiquity of AI.

Part I: Defending America in the AI Era.

Technology so ubiquitous in other facets of society will have an equivalent impact on international competition and conflict.⁴ We must adopt AI to change the way we defend America, deter adversaries, use intelligence to make sense of the world, and fight and win wars. The men and women who protect the United States must be able to leverage the AI and associated technologies that can help them accomplish their missions as quickly and safely as possible.

AI is the quintessential “dual use” technology—it can be used for civilian and military purposes. The AI promise—that a machine can perceive, decide, and act more quickly, in a more complex environment, with more accuracy than a human—represents a competitive advantage in any field. It will be employed for military ends, by governments and non-state groups.

We can expect the large-scale proliferation of AI-enabled capabilities. Many national security applications of AI will require only modest resources and good, but not great, expertise to use. AI algorithms are often accessible. The hardware is “off-the-shelf” and in most cases generally available to consumers (as with graphics processing units, for example). “Deepfake” capabilities can be easily downloaded and used by anyone.⁵ AI-enabled tools and mutating malware are in the hands of hackers.⁶ Cheap, lethal drones will be common. Azerbaijan’s use of Turkish drones and Israeli loitering munitions in combat against Armenia in October 2020 confirmed that autonomous military capabilities are spreading.⁷ Many states are watching and learning from these experiences. The likelihood of reckless or unethical uses of AI-enabled technologies by rogue states, criminals, or terrorists is increasing.

AI-enabled capabilities will be tools of first resort in a new era of conflict. State and non-state actors determined to challenge the United States, but avoid direct military confrontation, will use AI to amplify existing tools and develop new ones. Adversaries are exploiting our digital openness through AI-accelerated information operations and cyber attacks. Ad-tech will become natsec-tech as adversaries recognize what advertising and technology firms have recognized for years: that machine learning is a powerful tool for harvesting and analyzing data and targeting activities. Using espionage and publicly available data, adversaries will gather information and use AI to identify vulnerabilities in individuals, society, and critical infrastructure. They will model how best to manipulate behavior, and then act.

AI will transform all aspects of military affairs. AI applications will help militaries prepare, sense and understand, decide, and execute faster and more efficiently. Numerous weapon systems will leverage one or more AI technologies. AI systems will generate options for commanders and create battle networks connecting systems across all domains. It will transform logistics, procurement, training, and the design and development of new hardware. Adopting AI will demand the development of new tactics and operational concepts. In the future, warfare will pit algorithm against algorithm. The sources of battlefield advantage will shift from traditional factors like force size and levels of armaments to factors like superior data collection and assimilation, connectivity, computing power, algorithms, and system security.

Competitors are actively developing AI concepts and technologies for military use. Russia has plans to automate a substantial portion of its military systems.⁸ It has irresponsibly deployed autonomous systems in Syria for testing on the battlefield.⁹ China sees AI as the


path to offset U.S. conventional military superiority by “leapfrogging” to a new generation of technology. Its military has embraced “intelligentized war”—investing, for example, in swarming drones to contest U.S. naval supremacy.¹⁰ China's military leaders talk openly about using AI systems for “reconnaissance, electromagnetic countermeasures and coordinated firepower strikes.”¹¹ China is testing and training AI algorithms in military games designed around real-world scenarios. As these authoritarian states field new AI-enabled military systems, we are concerned that they will not be constrained by the same rigorous testing and ethical code that guide the U.S. military.

AI will revolutionize the practice of intelligence. There may be no national security function better suited for AI adoption than intelligence tradecraft and analysis. Machines will sift troves of data amassed from all sources, locate critical information, translate languages, fuse data sets from different domains, identify correlations and connections, redirect assets, and inform analysts and decision-makers. To protect the American people, perhaps the most urgent and compelling reason to accelerate the use of AI for national security is the possibility that more advanced machine analysis could find and connect the dots before the next attack, when human analysis alone may not see the full picture as clearly.

Defending against AI-capable adversaries without employing AI is an invitation to disaster. AI will compress decision time frames from minutes to seconds, expand the scale of attacks, and demand responses that will tax the limits of human cognition. Human operators will not be able to defend against AI-enabled cyber or disinformation attacks, drone swarms, or missile attacks without the assistance of AI-enabled machines. The best human operator cannot defend against multiple machines making thousands of maneuvers per second potentially moving at hypersonic speeds and orchestrated by AI across domains. Humans cannot be everywhere at once, but software can.

Compelling logic dictates quick, but careful and responsible, AI adoption. The government should adopt AI following the principle of legendary basketball coach John Wooden: “Be quick, but don't hurry.”¹² Like other “safety critical” applications of AI, military and intelligence functions require deliberation and caution before they are developed and fielded. Some current AI systems are narrow and brittle. All require rigorous testing, safeguards, and an understanding of how they might operate differently in the real world than in a testbed. AI-enabled autonomous weapon systems could be more precise, and as a result, reduce inadvertent civilian casualties. But they also raise important ethical questions about the role of human judgment in employing lethal force. If improperly designed or used, they could also increase the risk of military escalation.

*There is an emerging consensus on principles for using AI responsibly in the defense and intelligence communities.*¹³ If an AI-powered machine does not work as designed with predictability and guided by clear principles, then operators will not use it, organizations will not embrace it, and the American people will not support it. Hurrying would be counterproductive and dangerous if it caused Americans to lose confidence in the



“The best human operator cannot defend against multiple machines making thousands of maneuvers per second potentially moving at hypersonic speeds and orchestrated by AI across domains. Humans cannot be everywhere at once, but software can.”

benefits AI could confer. Risk, however, is inescapable. Failing to use AI to solve real national security challenges risks putting the United States at a disadvantage, leaving American service members more vulnerable, and spending taxpayer money unwisely on antiquated and inefficient equipment. Delaying AI adoption will push all of the risk onto the next generation of Americans—who will have to defend against, and perhaps fight, a 21st century adversary with 20th century tools.

The U.S. government still operates at human speed, not machine speed. Adopting AI requires profound adjustments in national security business practices, organizational cultures, and mindsets from the tactical to the strategic levels—from the battlefield to the Pentagon. The government lags behind the commercial state of the art in most AI categories, including basic business automation. It suffers from technical deficits that range from digital workforce shortages to inadequate acquisition policies, insufficient network architecture, and weak data practices. Bureaucracy is thwarting better partnerships with the AI leaders in the private sector that could help. The government must become a better customer and a better partner. National security innovation, in the absence of an impetus like a major war or terrorist attack, will require strong leadership.

Part II: Winning the Technology Competition.

In addition to AI's narrow national security and defense applications, AI is the fulcrum of a broader technology competition in the world. AI will be leveraged to advance all dimensions

of national power, from healthcare to food production to environmental sustainability. The successful adoption of AI in adjacent fields and technologies will drive economies, shape societies, and determine which states exert influence and exercise power in the world. Many countries have national AI strategies, but only the United States and China have the resources, commercial might, talent pool, and innovation ecosystem to lead the world in AI. In some areas of research and applications, China is already an AI peer, and it is more technically advanced in some applications.¹⁴ Within the next decade, China could surpass the United States as the world's AI superpower.¹⁵

On a level playing field, the United States is capable of out-innovating any competitor. However, today, there is a fundamental difference in the U.S. and China's approaches to AI innovation that puts American AI leadership in peril. For decades, the U.S. innovation model has been the envy of the world. The open exchange of ideas, free markets, and limited government involvement to support basic research are pillars of the American way of innovation and reflect American values. In America, tech firms compete for market share. They are not instruments of state power. Researchers collaborate in an open research environment in competition with their peers to make AI breakthroughs without regard for borders. The international flow of venture capital and AI-related commerce is encouraged as firms compete for profits and the next big idea.

Most AI progress in the United States should remain with the private sector and universities. We must not lose an innovation culture that is bottom-up and infused with a garage-startup mentality. However, a fully distributed approach is not a winning strategy in this strategic competition. Even large tech firms cannot be expected to compete with the resources of China or make the big investments the U.S. will need to stay ahead. We will need a hybrid approach meshing government and private-sector efforts to win the technology competition.

China is organized, resourced, and determined to win the technology competition. AI is central to China's global expansion, economic and military power, and domestic stability. It has a head start on executing a national AI plan as part of larger plans to lead the world in several critical and emerging technology fields. Beginning in 2017, China established AI goals, objectives, and strategies tied to specific timelines with resources backed by committed leadership to lead the world in AI by 2030.¹⁶ China is executing a centrally directed systematic plan to extract AI knowledge from abroad through espionage, talent recruitment, technology transfer, and investments. It has ambitious plans to build and train a new generation of AI engineers in new AI hubs. It supports "national champion" firms (including Huawei, Baidu, Alibaba, Tencent, iFlytek, and SenseTime) to lead development of AI technologies at home, advance state-directed priorities that feed military and security programs under the rubric of military-civil fusion, and capture markets abroad.¹⁷ It funds massive digital infrastructure projects across several continents. China developed an intellectual property (IP) strategy and is trying to set global technical standards for AI development.¹⁸ And its laws make it all but impossible for a company in China to shield its data from the authorities.¹⁹

Advancements in AI are contributing to a broad platform technology competition in e-commerce, search engines, social media, and much else. The countries, companies, and researchers that win the AI competition—in computing, data, talent, and commercialization—will be positioned to win a much larger game. In essence, more and better data, fed by a larger consumer/participant base, produce better algorithms, which produce better results, which in turn produces more users, more data, and better performance—until, ultimately, fewer companies will become entrenched as the dominant platforms. If China’s firms win these competitions, it will not only disadvantage U.S. commercial firms, it will also create the digital foundation for a geopolitical challenge to the United States and its allies. Platform domination abroad allows China to harvest the data of its users and permits China to extend aspects of its domestic system of control. Wherever China controls the digital infrastructure, social media platforms, and e-commerce, it would possess greater leverage and power to coerce, propagandize, and shape the world to conform to its goals.

The AI competition is complicated by deep interconnections. The United States and China are not operating in parallel lanes like the Soviets and Americans did in the space race, with disconnected research and development (R&D) enterprises and minimal commercial contacts. The research ecosystems in China and the United States are deeply connected through shared research projects, talent circulation (particularly from China to the United States), and commercial linkages that include supply chains, markets, and joint research ventures. It would be counterproductive to sever the technology ties to China that benefit basic research and U.S. companies. However, the United States must protect the integrity of open research, prevent the theft of American IP, and employ targeted tools like export controls and investment screening to protect technology industries critical to national security.

The United States retains advantages in critical areas, but trends are concerning. The world’s best scientific talent is more likely to stay home or migrate elsewhere today than in our recent past.²⁰ The U.S. lead in microelectronics—the hardware on which all AI runs—has diminished, and for cutting-edge chips it is dependent on foreign supply chains and manufacturers in Asia that are vulnerable to coercion or disruption.²¹ While many machine learning tools are widely available and per-unit computing costs have declined, the computing power and data access needed for cutting-edge deep learning research breakthroughs are making it harder for university-based researchers and smaller companies to compete.²² The geography of innovation remains concentrated in only some parts of the country.²³

The U.S. government must take a hands-on approach to national technology competitiveness. Promoting a diverse and resilient R&D ecosystem and commercial sector is a government responsibility. Expanding talent pipelines to attract the world’s best and redoubling efforts to educate AI-ready Americans are public policy choices. Judiciously, but aggressively, protecting critical AI intellectual property and thwarting the systemic

campaign of illicit knowledge transfer being conducted by competitors is a government obligation. Protecting hardware advantages and building resiliency into supply chains necessitate legislation and federal incentives. Bringing together like-minded allies and partners to build an international coalition that ensures a democratic vision for AI that will shape the digital future requires U.S.-led diplomacy.

The AI competition will require White House leadership. The critical elements of the strategy are too complicated for any one department or agency to lead because they cut across national security, economic, and technology policy. Only strong executive leadership from the White House can drive policy, force tradeoffs, and mobilize the country to make the necessary investments.

AI for What Ends? Technology and Values.

The widespread adoption of AI by governments around the world is impacting not only the international order among states, but also the political order within them. The stakes of the AI future are intimately connected to the enduring contest between authoritarian and democratic political systems and ideologies.

Technology itself does not possess an ideology, but how it is designed, where it is employed, and which laws govern its use reflect the priorities and values of those who design and employ it. More AI-enabled surveillance and analysis capabilities will soon be in the hands of most or all governments. As the technology diffuses, the main difference between states will have less to do with the quality or sophistication of the technology and more to do with the way it is used—for what purpose, and under what rules.

Authoritarian regimes will continue to use AI-powered face recognition, biometrics, predictive analytics, and data fusion as instruments of surveillance, influence, and political control. China's use of AI-powered surveillance technologies to repress its Uyghur minority and monitor all of its citizens foreshadows how authoritarian regimes will use AI systems to facilitate censorship, track the physical movements and digital activities of their citizens, and stifle dissent.²⁴ The global circulation of these digital systems creates the prospect of a wider adoption of authoritarian governance. But liberal democracies also employ AI for internal security and public safety purposes. More than half of the world's advanced democracies use AI-enabled surveillance systems.²⁵ Such technologies have legitimate public purposes and are compatible with the rule of law. Yet in states edging toward illiberal practices, utilizing digital tools in ways that undermine the rule of law could tip the scales toward further democratic backsliding. The preservation of individual liberties calls for continued vigilance. A responsible democracy must ensure that the use of AI by the government is limited by wise restraints to comport with the rights and liberties that define a free and open society.

The U.S. government should develop and field AI-enabled technologies with adequate transparency, strong oversight, and accountability to protect against misuse. Merely stating U.S.

opposition to the authoritarian use of AI is not enough. The United States must also demonstrate how a democracy should use AI to protect the security of its citizens in ways that uphold liberal democratic values. There is an urgent need to field AI for national security purposes against, for instance, foreign and domestic terrorists operating within our borders. There is also an enduring need to ensure that security applications of AI conform to core values of individual liberty and equal protection under law.

The United States must lead a coalition of democracies. As we ensure that AI is developed and used in ways that are safe for democracy at home, we must also promote global norms to make its use safe for democracy abroad. While the U.S. government's ability to influence the governance practices of other states is limited, a strong plank of the U.S. foreign policy agenda with respect to AI must be to promote human rights and counter techno-authoritarian trends. The United States can use diplomacy and leverage its global partnerships to advocate for establishing privacy-protecting technical standards and norms in international bodies, and it can work with like-minded nations to ensure that other nations have an alternative to embracing China's technology and methods of social control and access to technologies that protect democratic values like privacy. We do not seek a fragmented digital world. We want the United States and its allies to exist in a world with a diverse set of choices in digital infrastructure, e-commerce, and social media that will not be vulnerable to authoritarian coercion and that support free speech, individual rights, privacy, and tolerance for differing views.

Conclusion

We are at the beginning of the beginning of this new era of competition. We now know the uses of AI in all aspects of life will grow and the pace of innovation will accelerate. We know adversaries are determined to turn AI capabilities against us. We know a competitor is determined to surpass us in AI leadership. We know AI is accelerating breakthroughs in a wide array of fields. We know that whoever translates AI developments into applications first will have the advantage. Now we must act. The principles we establish, the federal investments we make, the national security applications we field, the organizations we redesign, the partnerships we forge, the coalitions we build, and the talent we cultivate will set America's strategic course. The United States should invest what it takes to maintain its innovation leadership, to responsibly use AI to defend free people and free societies, and to advance the frontiers of science for the benefit of all humanity. AI is going to reorganize the world. America must lead the charge.

Introduction - Endnotes

- ¹ *The Google DeepMind Challenge Match*, DeepMind (last accessed Jan. 7, 2021), <https://deepmind.com/alphago-korea>.
- ² Quoted in Orison Swett Marden, *How They Succeeded: Life Stories of Successful Men Told by Themselves*, Lothrop Publishing Co. at 238 (1901).
- ³ Andrew Ng is widely credited with making this comparison. See e.g., Shana Lynch, *Andrew Ng: Why AI Is the New Electricity*, Insights by Stanford Business (March 11, 2017), <https://www.gsb.stanford.edu/insights/andrew-ng-why-ai-new-electricity>.
- ⁴ For an overview of AI and international relations see Michael Horowitz, *Artificial Intelligence, International Competition, and the Balance of Power*, Texas National Security Review (May 2018), <https://doi.org/10.15781/T2639KP49>.
- ⁵ Karen Hao & Will Douglas Heaven, *The Year Deepfakes Went Mainstream*, MIT Technology Review (Dec. 24, 2020), <https://www.technologyreview.com/2020/12/24/1015380/best-ai-deepfakes-of-2020/>.
- ⁶ Nicholas Duran, et al., *2018 Webroot Threat Report*, Webroot (2018), https://www-cdn.webroot.com/9315/2354/6488/2018-Webroot-Threat-Report_US-ONLINE.pdf; *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, & Medicine (2019), <https://www.nap.edu/catalog/25488/implications-of-artificial-intelligence-for-cybersecurity-proceedings-of-a-workshop>; Ben Buchanan, et al., *Automating Cyber Attacks: Hype and Reality*, Center for Security and Emerging Technology (Nov. 2020), <https://cset.georgetown.edu/research/automating-cyber-attacks/>; *Deep Exploit: Fully Automatic Penetration Test Tool Using Deep Reinforcement Learning*, GitHub (last accessed Jan. 9, 2021), https://github.com/130-bbr-bbq/machine_learning_security/tree/master/DeepExploit.
- ⁷ Robyn Dixon, *Azerbaijan's Drones Owned the Battlefield in Nagorno-Karabakh—and Showed Future of Warfare*, Washington Post (Nov. 11, 2020), https://www.washingtonpost.com/world/europe/nagorno-karabakh-drones-azerbaijan-aremenia/2020/11/11/441bcbcd2-193d-11eb-8bda-814ca56e138b_story.html.
- ⁸ Vadim Kozyulin, *Militarization of AI*, Stanley Center for Peace and Security (July 2019), <https://stanleycenter.org/wp-content/uploads/2020/05/MilitarizationofAI-Russia.pdf>.
- ⁹ Dylan Malyasov, *Combat Tests in Syria Brought to Light Deficiencies of Russian Unmanned Mini-tank*, Defence Blog (June 18, 2018), <https://defence-blog.com/news/army/combat-tests-syria-brought-light-deficiencies-russian-unmanned-mini-tank.html>.
- ¹⁰ Testimony of Elsa Kania before the U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion* (June 7, 2019), https://www.uscc.gov/sites/default/files/June%207%20Hearing_Panel%201_Elsa%20Kania_Chinese%20Military%20Innovation%20in%20Artificial%20Intelligence_0.pdf; Elsa Kania, "AI Weapons" in China's Military Innovation, Brookings at 1 (April 20, 2020), https://www.brookings.edu/wp-content/uploads/2020/04/FP_20200427_ai_weapons_kania_v2.pdf.
- ¹¹ Marcus Clay, *The PLA's AI Competitions*, The Diplomat (Nov. 5, 2020), <https://thediplomat.com/2020/11/the-plas-ai-competitions/>.
- ¹² Andrew Hill & John Wooden, *Be Quick—But Don't Hurry: Finding Success in the Teachings of a Lifetime*, Simon & Schuster at 69 (2001).
- ¹³ Press Release, U.S. Department of Defense, *DoD Adopts Ethical Principles for Artificial Intelligence* (Feb. 24, 2020), <https://www.defense.gov/Newsroom/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>; *Principles of Artificial Intelligence Ethics for the Intelligence Community*, ODNI (last accessed Jan. 11, 2021), <https://www.dni.gov/index.php/features/2763-principles-of-artificial-intelligence-ethics-for-the-intelligence-community>.
- ¹⁴ Graham Allison & Eric Schmidt, *Is China Beating the U.S. to AI Supremacy?*, Belfer Center for Science and International Affairs (Aug. 2020), <https://www.belfercenter.org/publication/china-beating-us-ai-supremacy>.
- ¹⁵ See e.g., Alexandra Mousavizadeh, et al., *The Global AI Index*, Tortoise Media (Dec. 3, 2019), <https://www.tortoisemedia.com/2019/12/03/global-ai-index/>.

¹⁶ See Graham Webster, et al., *Full Translation: China's 'New Generation Artificial Intelligence Development Plan'*, New America (Aug. 1, 2017), <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/> (translating China's State Council Notice on the Issuance of the New Generation Artificial Intelligence Development Plan, dated July 20, 2017).

¹⁷ Benjamin Larsen, *Drafting China's National AI Team for Governance*, New America (Nov. 18, 2019), <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/drafting-chinas-national-ai-team-governance/> (originally published in Graham Webster, ed., *AI Policy and China: Realities of State-Led Development*, Stanford-New America DigiChina Project (Oct. 29, 2019), <https://d1y8sb8igg2f8e.cloudfront.net/documents/DigiChina-AI-report-20191029.pdf>); Meng Jing, *China to Boost Its 'National Team' to Meet Goal of Global AI Leadership by 2030*, South China Morning Post (Nov. 15, 2018), <https://www.scmp.com/tech/innovation/article/2173345/china-boost-its-national-team-meet-goal-global-ai-leadership-2030>; Gregory C. Allen, *Understanding China's AI Strategy*, Center for a New American Security (Feb. 6, 2019), <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy> ("The price of SenseTime and the other AI Champions being allowed to dominate these technologies is the Champions' extensive cooperation with China's national security community. Even beyond direct cooperation, China's success in commercial AI and semiconductor markets brings funding, talent, and economies of scale that both reduce China's vulnerability from losing access to international markets and offer useful technology for the development of weaponry and espionage capabilities.").

¹⁸ Emily de La Bruyère & Nathan Picarsic, *China Standards 2035*, Horizon Advisory (April 2020), <https://www.horizonadvisory.org/china-standards-2035-first-report>.

¹⁹ Murray Scot Tanner, *Beijing's New National Intelligence Law: From Defense to Offense*, Lawfare (July 20, 2017), <https://www.lawfareblog.com/beijings-new-national-intelligence-law-defense-offense>.

²⁰ Remco Zwetsloot, *China's Approach to Tech Talent Competition: Policies, Results, and the Developing Global Response*, Brookings Institution (April 2020), https://www.brookings.edu/wp-content/uploads/2020/04/FP_20200427_china_talent_policy_zwetsloot.pdf.

²¹ "Ninety percent of all high-volume, leading-edge [semiconductor] production will soon be based in Taiwan, China, and South Korea." The Department of Defense estimates that by 2022 only 8% of all semiconductor fabrication will occur in the United States, "down from 40% in the 1990s." Michaela D. Platzer, et al., *Semiconductors: U.S. Industry, Global Competition, and Federal Policy*, Congressional Research Service at 12 (Oct. 26, 2020), <https://fas.org/sqp/crs/misc/R46581.pdf> (quoting Rick Switzer, *U.S. National Security Implications of Microelectronics Supply Chain Concentration in Taiwan, South Korea, and the People's Republic of China*, U.S. Air Force [Sept. 2019]).

²² For example, non-elite universities and AI startups have difficulty affording the cost of compute resources and data for training sophisticated machine learning (ML) models. Nur Ahmed & Muntasir Wahed, *The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, arXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>. The need to bolster the nationwide AI infrastructure is the first recommendation in the 20-Year Roadmap issued by the Association for the Advancement of Artificial Intelligence. See *A 20-Year Community Roadmap for Artificial Intelligence Research in the US*, Computing Community Consortium and AAAI, at 3 (August 2019), <https://cra.org/ccc/wp-content/uploads/sites/2/2019/08/Community-Roadmap-for-AI-Research.pdf>.

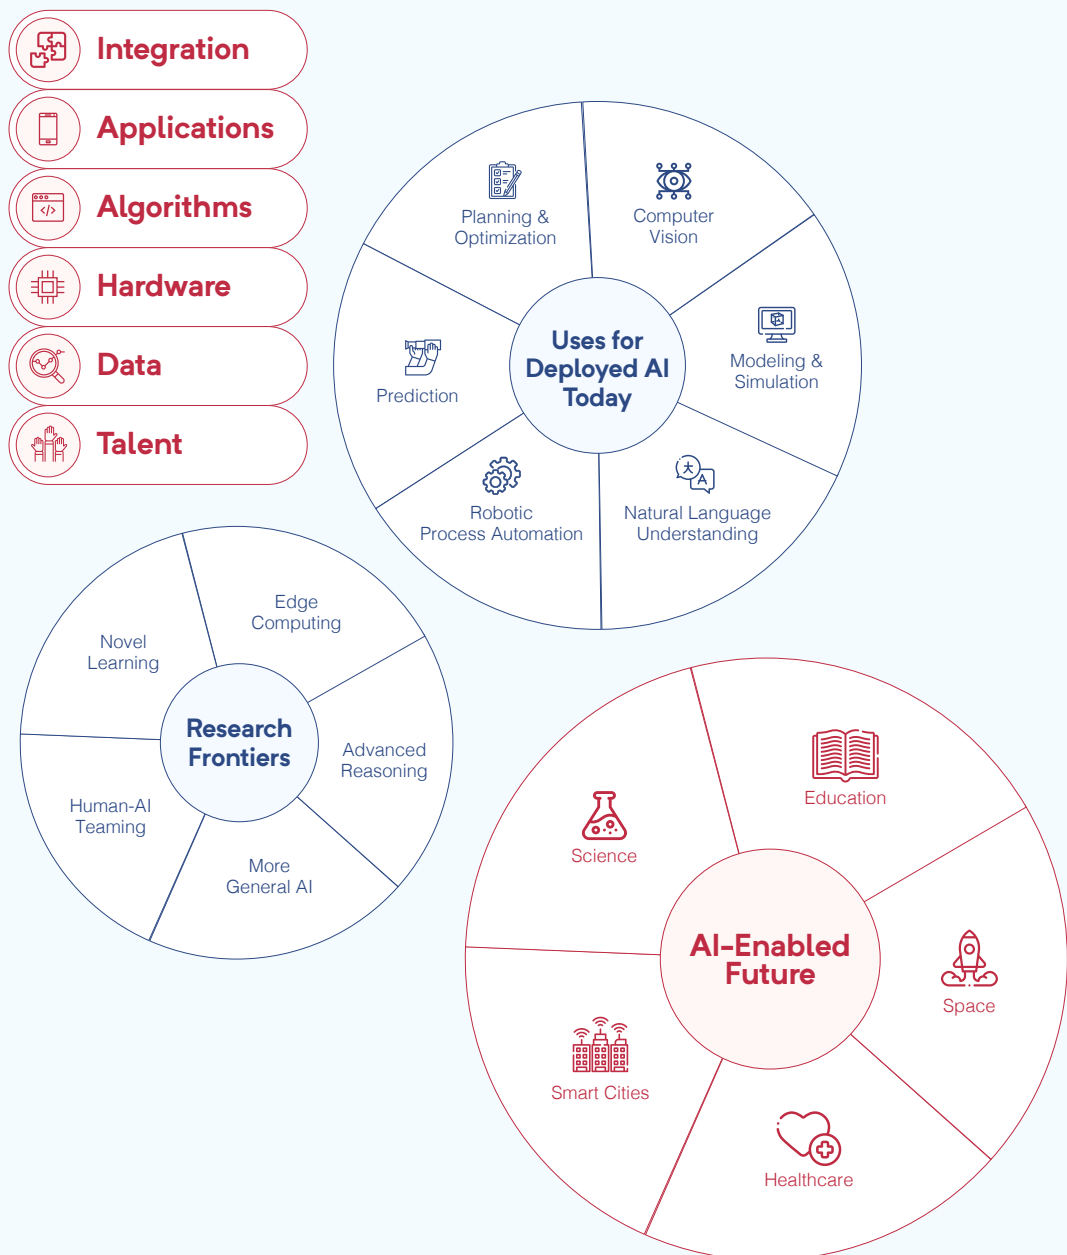
²³ More than 90% of U.S. innovation-sector job creation occurred in just five major coastal cities between 2005 and 2017. Robert D. Atkinson, et al., *The Case for Growth Centers: How to Spread Tech Innovation Across America*, Brookings (Dec. 9, 2019), <https://www.brookings.edu/research/growth-centers-how-to-spread-tech-innovation-across-america/>.

²⁴ See, e.g., Patrice Taddonio, *How China's Government Is Using AI on Its Uighur Muslim Population*, PBS Frontline (Nov. 21, 2019), <https://www.pbs.org/wgbh/frontline/article/how-chinas-government-is-using-ai-on-its-uighur-muslim-population/>; Bethany Allen-Ebrahimian, *Exposed: China's Operating Manuals for Mass Internment and Arrest by Algorithm*, International Consortium of Investigative Journalists (Nov. 24, 2019), <https://www.icij.org/investigations/china-cables/exposed-chinas-operating-manuals-for-mass-internment-and-arrest-by-algorithm/>.

²⁵ See Steven Feldstein, *The Global Expansion of AI Surveillance*, Carnegie Endowment for International Peace (Sept. 2019), <https://carnegieendowment.org/2019/09/17/global-expansion-of-ai-surveillance-pub-79847>.

Artificial Intelligence in Context

Artificial intelligence (AI) is not a single piece of hardware or software, but rather a constellation of technologies that depend on interrelated elements that can be envisioned as a stack.



Artificial Intelligence (AI) is not a single piece of hardware or software, but rather a constellation of technologies. To address such a broad topic, the Commission’s legislative mandate provided guidance on how to scope its work to include technologies that solve tasks requiring human-like perception, cognition, planning, learning, communication, or physical action; and technologies that may learn and act autonomously, whether in the form of software agents or embodied robots.¹

Successful development and fielding of AI technologies depends on a number of interrelated elements that can be envisioned as a stack.² AI requires talent, data, hardware, algorithms, applications, and integration. We regard talent as the most essential requirement because it drives the creation and management of all the other elements. Data is critical for most AI systems.³ Labeled and curated data enables much of current machine learning (ML) used to create new applications and improve the performance of existing AI applications. The underlying hardware provides the computing power to analyze ever-growing data pools and run applications. This hardware layer includes cloud-based compute and storage, supported by a networking and communications backbone, instrumental for connecting smart sensors and devices at the network edge. Algorithms are the mathematical operations that tell the system how to navigate the data to provide answers in response to specific questions. An application makes the answers useful for specific tasks. Integration of these elements is critical to fielding a successful end-to-end AI system. This requires significant engineering talent and investment to integrate existing data flows, decision pipelines, legacy equipment, testing designs, etc. This task of integration can be daunting and historically has been underestimated.⁴

AI technologies and applications such as pattern recognition, ML, computer vision, natural language understanding, and speech recognition have evolved for many decades. In the early years of AI, the period the Defense Advanced Research Projects Agency (DARPA) describes as the “first wave,” researchers explored many approaches, including symbolic logic, expert systems, and planning. Some of the most effective results were based on “handcrafted knowledge” defined by humans and then used by the machine for reasoning and interacting.⁵

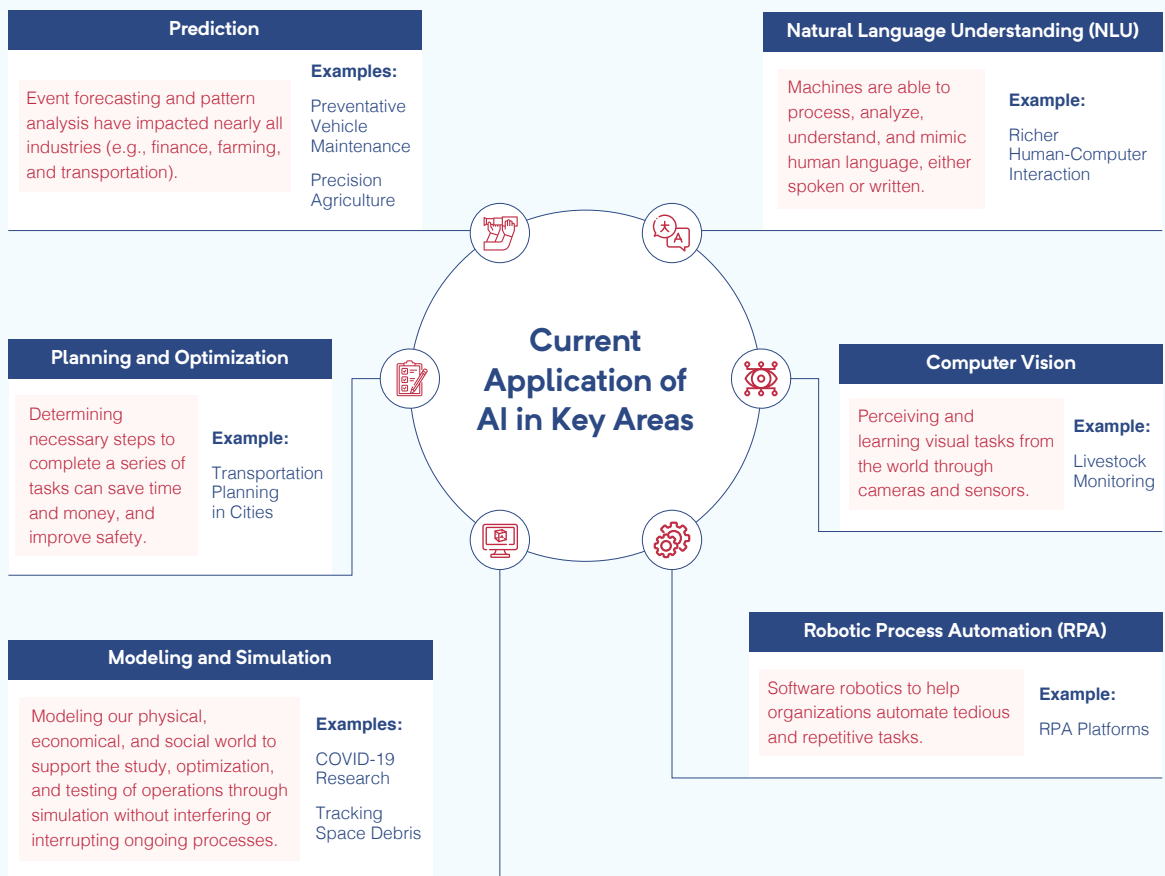
Within the past 10 years, we have witnessed a “second wave” of AI, propelled by large-scale statistical ML that enables engineers to create models that can be trained to specific problem domains if given exemplar data or simulated interactions. Learning from data, these systems are designed to solve specific tasks and achieve particular goals with competencies that, in some respects, parallel the cognitive processes of humans: perceiving, reasoning, learning, communicating, deciding, and acting. Today most fielded large-scale AI systems employ elements of both first- and second-wave AI approaches.

Age of Deployed AI.

Today, we have reached an inflection point. Global digital transformation has led to an overwhelming supply of data. Statistical ML algorithms, particularly deep neural networks, have matured as problem solvers—albeit with limitations.⁶ The powerful and networked computing that fuels ML capabilities has become widely available. The convergence of these factors now places this capable technology in the hands of the technical and non-technical alike. The fundamental “question is no longer how this technology works, but what it can do for you.”⁷

While the current technology still has significant limitations, it is well-suited for certain use cases. We have entered the age of deployed AI. AI is now ubiquitous, embedded in devices we use and interact with on a daily basis—for example, in our smartphones, wireless routers, and cars. We routinely rely on AI-enriched applications, whether searching for a new restaurant, navigating traffic, selecting a movie, or getting customer service over the phone or online.

Current Application of AI in Key Areas.



Forecasting the future of AI is difficult. Five years ago, few would have predicted the recent breakthroughs in natural language understanding that have resulted in systems that can generate full text almost indistinguishable from human prose.⁸ With a remarkable increase

of investments in the global AI industry over the past five years⁹ and an unprecedented amount of general R&D dollars being invested worldwide,¹⁰ there is no AI slowdown in sight—only new horizons for deployed AI.



“With a remarkable increase of investments in the global AI industry over the past five years and an unprecedented amount of general R&D dollars being invested worldwide, there is no AI slowdown in sight—only new horizons for deployed AI.”

Frontiers of AI Technology.

The next decade of AI research will likely be defined by efforts to incorporate existing knowledge, push forward novel ways of learning, and make systems more robust, generalizable, and trustworthy.¹¹ Research on advancing human-machine teaming will be at the forefront, as will improvements in hybrid AI techniques, enhanced training methods, and explainable AI.

Human-AI Teaming. Mastering human-AI collaboration and teaming is a foundational element for future application of AI. Synergy between humans and AI holds the promise of a whole greater than the sum of its parts. Researchers are addressing this challenge by studying issues of delegated authority, observability, predictability, directability, and trust.¹² Gaining greater understanding of how humans will learn to work with AI will provide insights for creating effective training programs for humans. Advances in language understanding are being pursued to create systems that can summarize complex inputs and engage through human-like conversation, a critical component of next-generation teaming. The frontier of teaming includes the need for collaborative intelligence among cohorts of agents, whether mixed groups of humans and machines or teams of coordinating machines.

Novel Ways of Learning. New learning methods are allowing for greater efficiency in both training and inference from data.¹³ This decreases dependence on vast data sets and widens the aperture of systems to handle tasks beyond their original scope, building pathways toward contextual learning and commonsense reasoning. Hybrid AI techniques combine different AI approaches to capitalize on their complementary strengths.¹⁴ For example, neuro-symbolic research is combining symbolic manipulation with neural networks.¹⁵ Model-based and data-based approaches may also be combined; for example, leveraging physics knowledge within statistical ML frameworks.¹⁶ Researchers are also advancing supervised learning techniques with low supplies of labeled data,¹⁷ while others have devised more efficient methods of labeling data.¹⁸ Synthetic data generation through simulation is one such promising approach.¹⁹ It allows a model to see conditions and scenarios it may not have encountered with a real data set, while preserving relationships between important variables in the original data and privacy of sensitive data.²⁰

Edge Computing. Breaking size, weight, and power barriers also increases the ubiquity of AI and aids privacy protection. Companies are working to pack more computational power into tighter, specialized chips that use less energy to train and run the same models. Such chips allow consumer devices to run complex models locally, rather than transmit data externally and wait for models to run remotely. Retaining data entirely on the device where a model is being trained or run is an advancement that could potentially enhance individual privacy in AI-powered systems.²¹

Advances in Reasoning. In comparison to humans, even our most capable current AI systems lack what one might think of as “commonsense reasoning.” Efforts are underway to create systems that can generalize knowledge and translate learning across domains. An AI system endowed with commonsense reasoning could effectively model the human ability to make and exploit presumptions about the physical properties, purpose, intentions, and behavior of people and objects and thereby characterize the probable consequences of an action or interaction. Advancements in categorization, in creating generalized structured ontologies, and in language understanding will drive the ability for machines to learn while understanding context and content and allow people to discover rapid solutions to problems that would historically take years to examine.²² This research promises to pave the way for more explainable AI along with greater ability to detect and mitigate bias, which will be essential to improving trustworthiness of these more general AI technologies.²³

Toward More General Artificial Intelligence. AI solutions to date have demonstrated narrow and deep competencies, but with fundamental distinction from capabilities demonstrated by humans. Humans perform tasks by learning without explicit supervised signals; they generalize skills required for one task and apply them to other tasks; and they accrue, manipulate, and reason with large amounts of commonsense knowledge. Some researchers have used the phrase “artificial general intelligence” (AGI) to refer to a goal of extending AI beyond narrow, vertical wedges of expertise. Debates have focused on

whether there might be specific breakthroughs that would lead to more general, human-like capabilities or whether the field will more likely continue to push more general AI along one or more dimensions of skills. No matter what the perspective, significant progress across the research areas mentioned in this section will be required to create more general AI systems.²⁴ If achieved, more general AI methods could have enormous benefits, but could also introduce new risks if safety challenges are not addressed. While breakthroughs are in no way guaranteed, the United States should continue to research systems with more human-like capabilities, accompanied by commensurate investments to ensure that those systems are safe and controllable.

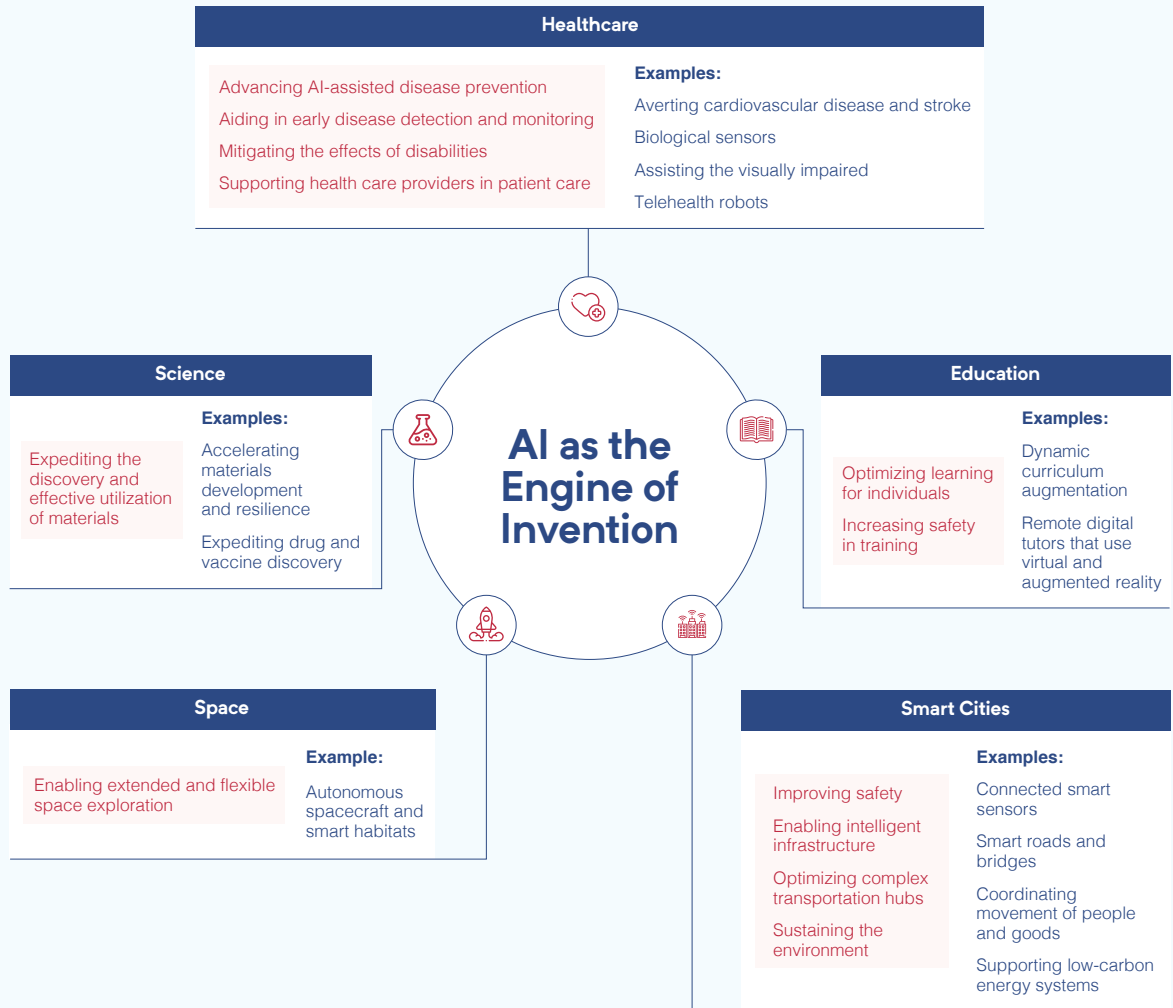
Advances in AI, including the mastery of more general AI capabilities along one or more dimensions, will likely provide new capabilities and applications. Some of these advances could lead to inflection points or leaps in capabilities. Such advances may also introduce new concerns and risks and the need for new policies, recommendations, and technical advances to assure that systems are aligned with goals and values,²⁵ including safety, robustness and trustworthiness.²⁶ The US should monitor advances in AI and make necessary investments in technology and give attention to policy so as to ensure that AI systems and their uses align with our goals and values.²⁷

Looking to an AI-Enabled Future.

Following the trajectories of the research threads outlined above sketches a future in which AI empowers humanity in unprecedented ways, unlocking capabilities across science, education, space technology, healthcare, infrastructure, manufacturing, agriculture, entertainment, and countless other sectors. For example, advances in natural language understanding could enable real-time, ubiquitous translation for more obscure languages for which written and spoken training data is limited.²⁸ This would transform the way we communicate across geographic and cultural barriers, enabling business, diplomacy, and free exchange of ideas.

Breakthroughs in integration of multi-modal, multi-source data could enable real-time AI-driven modeling and simulation for federal responses to crises including pandemics and natural disasters.²⁹ Drone feeds augmented with maps, building layouts, and other visual data layers could empower first responders with lifesaving emergency-scene understanding,³⁰ and AI could help build response plans, expedite command and control, and optimize logistics for a range of disaster-response scenarios.³¹

AI as the Engine of Invention.



Artificial Intelligence in Context - Endnotes

¹ The John S. McCain National Defense Authorization Act for Fiscal Year 2019 includes the following definition to guide the Commission's work: 1. Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve performance when exposed to data sets. 2. An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action. 3. An artificial system designed to think or act like a human, including cognitive architectures and neural networks. 4. A set of techniques, including machine learning that is designed to approximate a cognitive task. 5. An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision-making, and acting. See Pub. L. 115-232, 132 Stat. 1636, 1965 (2018).

² Andrew W. Moore, et al., *The AI Stack: A Blueprint for Developing and Deploying Artificial Intelligence*, Proc. SPIE 10635, Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR IX, 106350C (May 4, 2018), <https://doi.org/10.1117/12.2309483>; see also Dave Martinez, et al., *Artificial Intelligence: Short History, Present Developments, and Future Outlook*, MIT Lincoln Laboratory at 27 (Jan. 2019), <https://www.ll.mit.edu/media/9526>.

³ Note that model-based AI requires data for the manual construction of the model(s). Typically, this involves less data than statistical machine learning, but more human effort.

⁴ Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, ICSE-SEIP '19 Proceedings of the 41st International Conference on Software Engineering at 291-300 (2019), <https://2019.icse-conferences.org/details/icse-2019-Software-Engineering-in-Practice/30/Software-Engineering-for-Machine-Learning-A-Case-Study>; D. Sculley, et al., *Machine Learning: The High Interest Credit Card of Technical Debt*, SE4ML: Software Engineering for Machine Learning (NIPS 2014 Workshop), <https://ai.google/research/pubs/pub43146>.

⁵ John Launchbury, *A DARPA Perspective on Artificial Intelligence*, DARPA, 4-7 (Feb. 2017), <https://www.darpa.mil/attachments/AIFull.pdf>.

⁶ The limitations of today's statistical machine learning, as an example, include the vulnerability of unknowingly learning and amplifying biases in the training data; the fact that they are often complex models composed of a very large number of learned parameters, making them opaque and difficult to interpret; the fact that they are trained to solve narrow tasks and lack generalization to other related problems (such as when operationally encountered data fundamentally changes characteristic from the training data); and the fact that they require large amounts of labeled training data.

⁷ Andrew Moore, *When AI Becomes an Everyday Technology*, Harvard Business Review (June 7, 2019), <https://hbr.org/2019/06/when-ai-becomes-an-everyday-technology>.

⁸ Tom B. Brown, et al., *Language Models are Few-Shot Learners*, arXiv (July 22, 2020), <https://arxiv.org/abs/2005.14165>.

⁹ Zachary Arnold, et al., *Tracking AI Investment: Initial Findings from the Private Markets*, Center for Security and Emerging Technology (Sept. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Tracking-AI-Investment.pdf>.

¹⁰ According to UNESCO, global spending on R&D has reached a record high of almost US\$1.7 trillion. See *How Much Does Your Country Invest in R&D?*, UNESCO Institute for Statistics (last accessed Jan. 7, 2021), <http://uis.unesco.org/apps/visualisations/research-and-development-spending/>.

¹¹ For a recent debate from AI experts see *AI DEBATE 2: Moving AI Forward: An Interdisciplinary Approach*, Montreal Artificial Intelligence (Dec. 23, 2020), <https://montrealartificialintelligence.com/aidebate2.html>.

Artificial Intelligence in Context - Endnotes

¹² See e.g., Bryan Wilder, et al., *Learning to Complement Humans*, International Joint Conferences on Artificial Intelligence Organization (2020), <https://doi.org/10.24963/ijcai.2020/212>; Ece Kamar, et al., *Combining Human and Machine Intelligence in Large-scale Crowdsourcing*, Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012) (June 4-8, 2012), <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/galaxyZoo.pdf>; Ramya Ramakrishnan, et al., *Overcoming Blind Spots in the Real World: Leveraging Complementary Abilities for Joint Execution*, Proceedings of the AAAI Conference on Artificial Intelligence (July 17, 2019), <https://doi.org/10.1609/aaai.v33i01.33016137>; Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI '19: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (May 2019), <https://doi.org/10.1145/3290605.3300233>; Eric Horvitz, *Reflections on Challenges and Promises of Mixed-Initiative Interaction*, AI Magazine (June 15, 2007), <https://doi.org/10.1609/aimag.v28i2.2036>.

¹³ A goal of these new methods is to eliminate the need for many complex calculations that make traditional training very slow. Vincent Dutoridoir, *Sparse Gaussian Processes with Spherical Harmonic Features*, arXiv (June 30, 2020), <https://arxiv.org/abs/2006.16649>.

¹⁴ For a discussion on hybrid intelligence architectures that combine symbolic manipulation with deep learning see Gary Marcus, *The Next Decade in AI: Four Steps Towards Robust Artificial Intelligence*, arXiv at 14-19 (Feb. 19, 2020), <https://arxiv.org/abs/2002.06177>.

¹⁵ See *Neuro-symbolic AI*, MIT-IBM Watson AI Lab (last accessed Jan. 16, 2021), <https://mitibmwatsonailab.mit.edu/category/neuro-symbolic-ai/>.

¹⁶ Anuj Karpatne, et al., *Physics-guided Neural Networks (PGNN): An Application in Lake Temperature Modeling*, Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining (ACM SIGKDD) 2018 (Feb. 20, 2018), <https://arxiv.org/pdf/1710.11431.pdf>.

¹⁷ Rajat Raina, et al., *Self-Taught Learning: Transfer Learning from Unlabeled Data*, Proceedings of the 24th International Conference of Machine Learning (June 2007), <https://dl.acm.org/doi/abs/10.1145/1273496.1273592>; Dr. Bruce Draper, *Learning with Less Labeling*, DARPA (last accessed Dec. 19, 2020), <https://www.darpa.mil/program/learning-with-less-labeling>.

¹⁸ Rahul Dixit, et al., *Artificial Intelligence and Machine Learning in Sparse/Inaccurate Data Situations*, IEEE (Aug. 21, 2020), <https://ieeexplore.ieee.org/document/9172612>.

¹⁹ See Cem Dilmegani, *The Ultimate Guide to Synthetic Data in 2021*, AI Multiple (Jan. 12, 2021), <https://research.aimultiple.com/synthetic-data/>.

²⁰ *The Real Promise of Synthetic Data*, MIT News (Oct. 16, 2020), <https://news.mit.edu/2020/real-promise-synthetic-data-1016>.

²¹ *10 Breakthrough Technologies 2020*, MIT Technology Review (Feb. 26, 2020), <https://www.technologyreview.com/10-breakthrough-technologies/2020/#tiny-ai>.

²² *The World's Largest and Most Complete Common-Sense Knowledge Base*, Cycorp (last accessed Dec. 19, 2020), <https://www.cyc.com/the-cyc-platform/the-knowledge-base>; John Pavlus, *Common Sense Comes Closer to Computers*, Quanta Magazine (April 30, 2020), <https://www.quantamagazine.org/common-sense-comes-to-computers-20200430/>; Antoine Bosselut, et al., *COMET: Commonsense Transformers for Automatic Knowledge Graph Construction*, Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (2019), <https://homes.cs.washington.edu/~msap/pdfs/bosselut2019comet.pdf>.

²³ Amy Blumenthal, *How to Make AI Trustworthy*, Science Daily (Aug. 31, 2020), <https://www.sciencedaily.com/releases/2020/08/200827105937.htm>.

²⁴ See Benedict Neo, *Top 4 AI companies leading in the race towards Artificial General Intelligence*, Towards Data Science (April 13, 2020), <https://towardsdatascience.com/four-ai-companies-on-the-bleeding-edge-of-artificial-general-intelligence-b17227a0b64a>; Srishti Deoras, *9 Companies Doing Exceptional Work In AGI, Just Like OpenAI*, Analytics India Magazine (July 25, 2019), <https://analyticsindiamag.com/9-companies-doing-exceptional-work-in-agi-just-like-openai/>.

²⁵ In the *Key Considerations*, the Commission describes practices, technologies and operational policies to develop and field systems that align with key values. Importantly, agencies must consider values as (1) embodied in choices about engineering trade-offs and (2) explicitly represented in the goals and utility functions of an AI system. See *Key Considerations for Responsible Development & Fielding of Artificial Intelligence*, NSCAI (July 2020), <https://www.nscai.gov/wp-content/uploads/2021/01/Key-Considerations-for-Responsible-Development-Fielding-of-AI.pdf>. In addition to actions and investments needed now (see the Blueprint for Action associated with Chapter 7 of this report), the *Key Considerations* include policies and practices that should be updated to reflect new AI considerations as the technology evolves.

²⁶ This will require R&D as noted in the *Key Considerations* and Chapter 7 of this report. It will also require continued investment in system architectures to limit the consequences of system failure, to monitor AI performance as systems run to assess if they are performing as intended, and to overcome S&T gaps for audit and oversight. For more information, see Chapters 7 and 8 of this report.

²⁷ As noted in Chapter 8 and its Blueprint for Action, this will require, for instance, sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies. Disallowed outcomes and policy guidance will need to be updated over time as community norms and technical capabilities change. Further, the *Key Considerations* note that engineering practices will need to assess general feasibility and compliance with disallowed outcomes expressed in policy, as well as the demonstrated technical maturity of specific candidate AI technologies. See *Key Considerations for Responsible Development & Fielding of Artificial Intelligence*, NSCAI (July 2020), <https://www.nscai.gov/wp-content/uploads/2021/01/Key-Considerations-for-Responsible-Development-Fielding-of-AI.pdf>.

²⁸ Envisioned translation systems will leverage feedback to the system by actively correcting translation and recognition errors the software makes, improving performance as the interaction between translating parties goes on. A major goal is rapid deployment to new languages that have not been seen before. For example, Carnegie Mellon's DIPLOMAT Project makes interactive speech translation possible through a new architecture called Multi Engine Machine Translation (MEMT). DIPLOMAT gives users the ability to provide translation corrections to support rapid development to new languages that have not been seen before. Robert Frederking, *Interactive Speech Translation in the DIPLOMAT Project*, Carnegie Mellon University Language Technologies Institute (last accessed Dec. 19, 2020), <http://www.cs.cmu.edu/~air/papers/acl97-workshop.pdf>.

²⁹ Modeling and simulation can also help prepare for effective pandemic supply chain responses orchestrated by the government. Madhav Marathe, *High Performance Simulations to Support Real-time COVID19 Response*, SIGSIM-PADS '20 (June 2020), <https://dl.acm.org/doi/pdf/10.1145/3384441.3395993>.

³⁰ Edgybees (last accessed Dec. 19, 2020), <https://edgybees.com/>.

³¹ *Department of Energy Announces the First Five Consortium*, U.S. Department of Energy (Aug. 18, 2020), <https://www.energy.gov/articles/department-energy-announces-first-five-consortium>.

PART ONE



PART I: DEFENDING AMERICA IN THE AI ERA	41
Chapter 1: Emerging Threats in the AI Era	43
Chapter 2: Foundations of Future Defense	59
Chapter 3: AI and Warfare	75
Chapter 4: Autonomous Weapon Systems and Risks Associated with AI-Enabled Warfare	89
Chapter 5: AI and the Future of National Intelligence	107
Chapter 6: Technical Talent in Government	119
Chapter 7: Establishing Justified Confidence in AI Systems	131
Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security	141

Chapter 1: Emerging Threats in the AI Era

Societal Level of Conflict



AI-Enabled
Information
Operations



Data Harvesting
and Targeting of
Individuals



Accelerated
Cyber Attacks



Adversarial AI



AI-Enabled
Biotechnology

The U.S. government is not prepared to defend the United States in the coming artificial intelligence (AI) era. AI applications are transforming existing threats, creating new classes of threats, and further emboldening state and non-state adversaries to exploit vulnerabilities in our open society.¹ AI systems will extend the range and reach of adversaries into the United States just as the missile age and terrorism brought threats closer to home. Because of AI, adversaries will be able to act with micro-precision, but at macro-scale and with greater speed. They will use AI to enhance cyber attacks and digital disinformation campaigns and to target individuals in new ways. AI will also help create precisely engineered biological agents. And adversaries will manipulate the AI systems we will rely upon.

How AI is Transforming the Threat Landscape

Current Threats Advanced BY AI Systems

AI transforms existing range and reach of threats

- Self-replicating AI-generated malware
- Improved and autonomous disinformation campaigns
- AI-engineered and targeted pathogens

New Threats FROM AI Systems

AI creates new threat phenomena

- Deepfakes and computational propaganda
- Micro-targeting: AI-fused data for targeting or blackmail
- AI swarms and nano-swarms

Threats TO AI Stacks Themselves

AI itself is also a new attack surface

- AI attack involves the whole “AI stack”. Examples include:
 - Model inversion
 - Training data manipulation
 - “Data lake” poisoning

Future Threats VIA AI Systems

Examples of potential threats to keep in view


- Rapid machine-to-machine escalation via automated C2
- AI-enabled human augmentation by peer competitors
- Proliferation of simple lethal autonomous weapons to terrorists

AI technologies exacerbate two existing national security challenges:

- First, digital dependence in all walks of life increases vulnerabilities to cyber intrusion across every segment of our society: corporations, universities, government, private organizations, and the homes of individual citizens. In parallel, new sensors have flooded the modern world. The internet of things (IoT), cars, phones, homes, and social media platforms collect streams of data, which can then be fed into AI systems that can identify, target, and manipulate or coerce our citizens.²
- Second, state and non-state adversaries are challenging the United States below the threshold of direct military confrontation by using cyber attacks, espionage, psychological and political warfare, and financial instruments. Adversaries do not need AI to conduct widespread cyber attacks, exfiltrate troves of sensitive data about American citizens, interfere in our elections, or bombard us with malign information on


digital platforms. However, AI is starting to change these attacks in kind and in degree, creating new threats to the U.S. economy, critical infrastructure, and societal cohesion.³ Moreover, these AI-enabled capabilities will be used across the spectrum of conflict. They will be used as tools of first resort in non-military conflicts, as a prelude to military actions, or in concert with military actions in war.

Americans are waking to some of the privacy implications of their digital dependence and the potential threats from AI-powered malign information, like deep fakes. However, debate in the United States has not yet accounted for the full scope and danger of the AI-enabled threats and the overall security risks to the AI systems all around us. The prospect of adversaries using machine learning (ML), planning, and optimization to create systems to manipulate citizens' beliefs and behavior in undetectable ways is a gathering storm.⁴ Most concerning is the prospect that adversaries will use AI to create weapons of mass influence to use as leverage during future wars, in which every citizen and organization becomes a potential target.



“The prospect of adversaries using machine learning, planning, and optimization to create systems to manipulate citizens’ beliefs and behavior in undetectable ways is a gathering storm. Most concerning is the prospect that adversaries will use AI to create weapons of mass influence to use as leverage during future wars, in which every citizen and organization becomes a potential target.”

Societal Level Impact.



The rest of this chapter discusses five AI-related threats that already have been, or soon will be, developed and used against the United States.

1. AI-Enabled Information Operations.

AI and associated technologies will increase the magnitude, precision, and persistence of adversarial information operations. AI exacerbates the problem of malign information in three ways:

- **Message.** AI can produce original text-based content and manipulate images, audio, and video, including through generative adversarial network (GAN)-enabled and reinforcement learning (RL) deep fakes that will be very difficult to distinguish from authentic messages.

- **Audience.** AI can construct profiles of individuals' preferences, behaviors, and beliefs to target specific audiences with specific messages.
- **Medium.** AI can be embedded within platforms, such as through ranking algorithms, to proliferate malign information.

AI-enabled malign information campaigns will not just send one powerful message to 1 million people, like 20th century propaganda. They also will send a million individualized messages—configured on the basis of a detailed understanding of the targets' digital lives, emotional states, and social networks.⁵ Rival states are already using AI-powered malign information. For example, according to Taiwan authorities, China's government tested its AI-powered malign information capacities during the 2020 Taiwan elections.⁶ A National Basketball Association general manager was harassed on social media for supporting protesters in Hong Kong, in an effort that may have involved autonomous bots.⁷ Other techniques rely on AI-generated fake personas.⁸ The control and manipulation of digital information has become central to the Kremlin's strategy, including in efforts to undermine the integrity of the democratic process in the United States and elsewhere.⁹

In the United States, the private sector has taken the leading role in combating foreign malign information. Social media companies in particular have extensive operations to track and manage information on their platforms. But coordination between the government and the social media firms remains ad hoc. We need a more integrated public-private response to the problem of foreign-generated disinformation. Moreover, the government needs to devote greater attention and resources to the technical challenges of detection, attribution, and media authentication. The government should:

Create a Joint Interagency Task Force and Operations Center. Congress has authorized a Foreign Malign Influence Response Center to be established within the Office of the Director of National Intelligence (ODNI).¹⁰ The government should use this authority to create a technologically advanced, 24-hour task force and operations center to lead and integrate government efforts to counter foreign-sourced malign information. It would survey the landscape of relevant public and private actors, coordinate among them, and act in real time to counter foreign information campaigns. To expose, attribute, and respond effectively, the center must be equipped with modern AI-enabled digital tools and staff with specialized expertise.

Recommendation

Fund the Defense Advanced Research Projects Agency (DARPA) to coordinate multiple research programs to detect, attribute, and disrupt AI-enabled malign information campaigns and to authenticate the provenance of digital media. Additional funding would amplify ongoing DARPA research programs to detect synthetic media and expand its efforts into attributing and disrupting malign information campaigns.¹¹ However promising some of these detection technologies may prove to be individually, funding to develop alternative technologies to authenticate the provenance of the digital media will provide a more technologically robust means to prevent the impersonation of trusted sources of information.¹² DARPA should pursue these programs and help transition all of these

Recommendation

technologies and applications to government departments and agencies, in order to assist with detecting, attributing, and disrupting malign information campaigns in real time.

Recommendation

Create a task force to study the use of AI and complementary technologies, including the development and deployment of standards and technologies, for certifying content authenticity and provenance. The White House Office of Science and Technology Policy should take the lead in creating this task force. In response to the challenges of misinformation, efforts are underway to develop standards and pipelines aimed at certifying the authenticity and provenance of audiovisual content.¹³ These efforts make use of technologies, including encryption and fragile watermarking, to secure and track the expected transformations of content via production and transmission pipelines. These efforts offer the opportunity to mitigate malign information campaigns that seek to corrupt or spoof highly trusted sources of information across our digital ecosystem. This technology area is ripe for public-private partnership. Several private organizations are already forming to fight disinformation efforts in this realm.¹⁴

2. Data Harvesting and Targeting of Individuals.

“Potential adversaries will recognize what every advertiser and social media company knows: AI is a powerful targeting tool.”

Data security is a national security problem. “Ad-tech” has become “natsec-tech.” Potential adversaries will recognize what every advertiser and social media company knows: AI is a powerful targeting tool. Just as AI-powered analytics transformed the relationship between companies and consumers, now it is transforming the relationship between governments and individuals. The broad circulation of personal data drives commercial innovation but also creates vulnerabilities.¹⁵ We fear that adversaries’ systematic efforts to harvest data on U.S. companies, individuals, and the government is about more than traditional espionage.¹⁶ Adversaries will combine widely available commercial data with data acquired illicitly—as in the 2015 Office of Personnel Management hack—to track, manipulate, and coerce individuals.¹⁷ The reach of tools that China, for instance, uses

to monitor, control, and coerce its own citizens—big data analytics, surveillance, and propaganda—can be extended beyond its borders and directed at foreigners.¹⁸ Without adequate data protection, AI makes it harder for anyone to hide his or her financial situation, patterns of daily life, relationships, health, and even emotions. Personal and commercial vulnerabilities become national security weaknesses as adversaries map individuals, networks, and social fissures in society; predict responses to different stimuli; and model how best to manipulate behavior or cause harm. The rise and spread of these techniques represent a major counterintelligence challenge.¹⁹

For the government to treat the data of its citizens and businesses as a national security asset, substantial changes are required in the way we think about data security and in our policies and laws to strengthen it. We need to identify categories and combinations of personal and commercial data that are most sensitive. Early efforts to limit foreign adversaries' data harvesting—such as the government's decision to force a Chinese company to relinquish ownership of a popular dating application for fear of what a hostile adversary could do with sensitive private data²⁰—represent important initial steps. However, the government lacks a broad approach with clear policies, criteria, or authorities to confront this multifaceted problem. The government should:


Develop policies that treat data security as national security, including in these areas:

Recommendation

- **First, from a technical standpoint, the government must ensure that a security development lifecycle approach is in place for its own AI systems (including commercial systems it acquires),** which should include a focus on potential privacy attacks.²¹ Red teaming must include privacy expertise. Government databases should be federated and anonymized whenever possible, and personal data retained no longer than is necessary, in order to make it more difficult for adversaries to utilize information for malicious purposes.
- **Second, the government should ensure that data privacy and security are priority considerations** as part of larger efforts to strengthen foreign investment screening and supply chain intelligence and risk management.²²
- **Third, national efforts to legislate and regulate data protection and privacy must integrate national security considerations,** such as limiting the ability of hostile foreign actors to acquire sensitive data on Americans on the commercial market.²³

3. Accelerated Cyber Attacks.

Malware in the AI era will be able to mutate into thousands of different forms once it is lodged on a computer system. Such mutating polymorphic malware already accounts for more than 90% of malicious executable files.²⁴ Deep RL tools can already find vulnerabilities, conceal malware, and attack selectively.²⁵ While it is uncertain which methods will dominate, there is a clear path for U.S. adversaries to transform the effectiveness of cyber attack and espionage campaigns with an ensemble of new and old algorithmic means to automate, optimize, and inform attacks.²⁶ This goes beyond AI-enhanced malware. Machine learning has current and potential applications across all the phases of cyber attack campaigns



“Machine learning has current and potential applications across all the phases of cyber attack campaigns ...”

and will change the nature of cyber warfare and cyber crime.²⁷ The expanding application of existing AI cyber capabilities will make cyber attacks more precise and tailored, further accelerate and automate cyber warfare, enable stealthier and more persistent cyberweapons, and make cyber campaigns more effective on a larger scale.

U.S. defenses have proven incapable of handling even more elementary cyber challenges. Vulnerabilities remain open in outdated infrastructure and medical devices, while new vulnerabilities are proliferating in 5G networks, billions of IoT devices, and in software supply chains.²⁸ The multibillion-dollar global damage caused by Russia’s 2017 NotPetya attack concretely demonstrates the power of even basic automated malware, the risk tolerance of capable state actors, and the consequences of such capabilities proliferating.²⁹ Though defensive applications of AI bring the promise to improve our national cyber defenses, AI can’t defend inherently vulnerable digital infrastructure. To address the present threat, Congress must continue implementing the Cyberspace Solarium Commission’s recommendations.³⁰ With this foundation for cyber defense, the U.S. can prepare for expanding threats via testing and building the instrumented infrastructure required for AI-enabled cyber defenses, establishing better incentives for security, properly organizing to meet the challenge, and keeping attackers off balance. Pervasive cyber-enabled espionage and attacks on U.S. computer networks and critical infrastructure will continue—and will become more damaging with AI—unless urgent federal action is taken. The government should:

Recommendation

Develop and deploy AI-enabled defenses against cyber attacks. National security agencies need to acquire the sensors and instrumentation needed to train AI systems to detect and respond to threats on their networks. AI-enabled cyber defenses will also need large-scale, instrumented, and realistic testing, and they must be robust enough to withstand adversarial attacks. The defenses should be employed to expand machine speed information sharing, behavior-based anomaly detection, and malware mitigation across government networks. To capitalize on these capabilities, the government should accelerate the establishment of a Joint Cyber Planning and Operations Center, modeled after the National Counterterrorism Center.³¹ The Center would serve as a centralized cyber intelligence sharing and collaboration unit with multi-agency jurisdiction and authorities to investigate threats, proactively support defensive mitigations, and coordinate responses.

4. Adversarial AI.

AI systems represent a new target for attack. While we are on the front edge of this phenomenon, commercial firms and researchers have documented attacks that involve evasion, data poisoning, model replication, and exploiting traditional software flaws to deceive, manipulate, compromise, and render AI systems ineffective.³² This threat is related to, but distinct from, traditional cyber activities, because AI systems will be vulnerable to adversarial attacks from any domain where AI augments action—civilian or military.³³ Given the reliance of AI systems on large data sets and algorithms, even small manipulations of these data sets or algorithms can lead to consequential changes for how AI systems operate. The threat is not hypothetical: adversarial attacks are happening and already impacting commercial ML systems.³⁴ With rare exceptions, the idea of protecting AI systems has been an afterthought in engineering and fielding AI systems, with inadequate investment in research and development.³⁵ Only three of 28 organizations recently surveyed have “the right tools in place to secure their ML systems.”³⁶ There has not yet been a uniform effort to integrate AI assurance across the entire U.S. national security enterprise. To improve AI “assurance,” the government should:

Create a National AI Assurance Framework. All government agencies will need to develop and apply an adversarial ML threat framework to address how key AI systems could be attacked and should be defended. An analytical framework can help to categorize threats to government AI systems and assist analysts with detecting, responding to, and remediating threats and vulnerabilities.³⁷

Recommendation

Create dedicated red teams for adversarial testing. Red teams should assume an offensive posture, trying to break systems and make them violate rules for appropriate behavior. Because of the scarcity of required expertise and experience for AI red teams, DoD and ODNI should consider establishing government-wide communities of AI red-teaming capabilities that could be applied to multiple AI developments.³⁸

Recommendation

5. AI-Enabled Biotechnology.

Biology is now programmable. New technologies such as the gene editing tool CRISPR ushered in an era where humans are able to edit DNA. Combined with massive computing power and AI, innovations in biotechnology may provide novel solutions for mankind’s most vexing challenges, including in health, food production, and environmental sustainability. Like other powerful technologies, however, applications of biotechnology can have a dark side. The COVID-19 pandemic reminded the world of the dangers of a highly contagious pathogen. AI may enable a pathogen to be specifically engineered for lethality or to target a genetic profile—the ultimate range and reach weapon. Also, AI, when applied to biology, could optimize for the physiological enhancement of human beings, including intelligence and physical attributes. To the extent that brain waves can be represented as a machine vision challenge for AI, the mysteries of the brain may be unlocked and programmed.

Individuals, societies, and states will have different moral and ethical views and accept different degrees of risk in the name of progress, and U.S. competitors are comparatively likely to take more risk-tolerant actions and conform less rigidly to bioethical norms and standards. China understands the tremendous upside associated with leading the bio revolution. Massive genomic data sets at places like BGI Group (formerly known as the Beijing Genomics Institute), coupled with China's now-global genetic data collection platform and "all-of-nation" approach to AI, will make them a formidable competitor in the bio realm.³⁹ BGI may be serving, wittingly or unwittingly, as a global collection mechanism for Chinese government genetic databases, providing China with greater raw numbers and diversity of human genome samples as well as access to sensitive personal information about key individuals around the world.⁴⁰ The United States cannot afford to look back in 10 years and be "surprised" by the biotechnology equivalent of Huawei. Additionally, Russia's long-standing disregard for scientific norms and bioethical principles, demonstrated by its development and employment of novel nerve agents such as Novichok for assassination attempts and U.S. government concerns over Russia's compliance with the Biological Weapons Convention, could presage a willingness to utilize advanced biotechnology abilities for nefarious purposes.⁴¹ The government should:

Recommendation

Increase the profile of biosecurity and biotechnology issues within U.S. national security agencies. Given how AI will substantially increase the rate of technical advancement in biotechnology, the government should update the National Biodefense Strategy to include a wider vision of biological threats, such as human enhancement, exploitation of genetic data for malicious ends, and ways U.S. competitors could utilize biotechnology or biodata advantages for novel purposes. Additionally, U.S. officials should warn of the dangers associated with foreign actors obtaining personal genetic information, specifically highlighting concerns about the links between BGI and the Chinese government.⁴²

Chapter 1 - Endnotes

¹ A threat can be understood as an adversary capability paired with a vulnerability that can create a harmful consequence. See Terry L. Deibel, *Foreign Affairs Strategy: Logic for American Statecraft*, Cambridge University Press at 142-150 (2007). Threats can be graded further by their seriousness, likelihood, imminence, and tractability.

² Stuart A. Thompson & Charlie Warzel, *Twelve Million Phones, One Dataset, Zero Privacy*, New York Times (Dec. 19, 2019), <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>. For example, the internet of things (IoT) and AI-powered applications can turn your new robotic vacuum into a listening device. See Sriram Sami, et al., *Spying with Your Robot Vacuum Cleaner: Eavesdropping via Lidar Sensors*, Proceedings of the 18th Conference on Embedded Networked Sensor Systems (Nov. 2020), <https://dl.acm.org/doi/10.1145/3384419.3430781>.

³ This is in some ways analogous to what Cold War strategists called “counter-value targeting.” See Lawrence Freedman, *The Evolution of Nuclear Strategy*, Palgrave Macmillan Vol. 20 at 119-122 (1989). In the realm of nuclear strategy, this was also known as counter-city or counter-economy targeting.

⁴ Some observers have used the concept of “sharp power” to describe such efforts to wield influence in open societies. These uses of power are sharp “in the sense that [authoritarian states aim to] pierce, penetrate, or perforate the information environments in the targeted countries.” *Sharp Power: Rising Authoritarian Influence*, National Endowment for Democracy at 13 (Dec. 5, 2017), <https://www.ned.org/sharp-power-rising-authoritarian-influence-forum-report/>. See also Testimony of Dr. Eric Horvitz, Microsoft, before the U.S. Senate Committee on Commerce, Science, & Transportation, Subcommittee on Space, Science, & Competitiveness, *Hearing on the Dawn of Artificial Intelligence* at 13 (Nov. 30, 2016), http://erichorvitz.com/Senate_Testimony_Eric_Horvitz.pdf.

⁵ Some have characterized AI-driven information operations as “computational propaganda.” See Matt Chessen, *The MADCOM Future: How Artificial Intelligence Will Enhance Computational Propaganda, Reprogram Human Culture, and Threaten Democracy... and What Can Be Done About It*, Atlantic Council (Sept. 2017), https://www.atlanticcouncil.org/wp-content/uploads/2017/09/The_MADCOM_Future_RW_0926.pdf.

⁶ Philip Sherwell, *China Uses Taiwan for AI Target Practice to Influence Elections*, The Australian (Jan. 5, 2020), <https://www.theaustralian.com.au/world/the-times/china-uses-taiwan-for-ai-target-practice-to-influence-elections/news-story/57499d2650d4d359a3857688d416d1e5>.

⁷ Ben Cohen, et al., *How One Tweet Turned Pro-China Trolls Against the NBA*, Wall Street Journal (Oct. 16, 2019), <https://www.wsj.com/articles/how-one-tweet-turned-pro-china-trolls-against-the-nba-11571238943>. On automated bots, see, e.g., Sarah Kreps & Miles McCain, *Not Your Father's Bots*, Foreign Affairs (Aug. 2, 2019), <https://www.foreignaffairs.com/articles/2019-08-02/not-your-fathers-bots>.

⁸ James Vincent, *An Online Propaganda Campaign Used AI-Generated Headshots to Create Fake Journalists*, The Verge (July 7, 2020), <https://www.theverge.com/2020/7/7/21315861/ai-generated-headshots-profile-pictures-fake-journalists-daily-beast-investigation>.

⁹ For recent studies on technical aspects of Russia’s interference in the 2016 election, see Alexander Spangher, et al., *Characterizing Search-Engine Traffic to Internet Research Agency Web Properties*, Web Conference (2020), <https://www.microsoft.com/en-us/research/publication/characterizing-search-engine-traffic-to-internet-research-agency-web-properties/>; Ryan Boyd, et al., *Characterizing the Internet Research Agency’s Social Media Operations During the 2016 U.S. Presidential Election Using Linguistic Analyses*, PsyArXiv Preprints (2018), <https://psyarxiv.com/ajh2q/>. See also Alina Polyakova, *Weapons of the Weak: Russian and AI-driven Asymmetric Warfare*, Brookings Institution (Nov. 15, 2018), <https://www.brookings.edu/research/weapons-of-the-weak-russia-and-ai-driven-asymmetric-warfare/>.

¹⁰ See Pub. L. 116-92, National Defense Authorization Act for Fiscal Year 2020, 133 Stat. 1198, 2129 (2019).

¹¹ These include the Media Forensics (MediFor) and Semantic Forensics (SemaFor) programs. See Dr. Matt Turek, *Media Forensics*, DARPA (last accessed Jan.10, 2021), <https://www.darpa.mil/program/media-forensics>; Dr. Matt Turek, *Semantic Forensics*, DARPA (last accessed Jan.10, 2021), <https://www.darpa.mil/program/semantic-forensics>.

Chapter 1 - Endnotes

¹² See, e.g., Paul England, et al., *AMP: Authentication of Media via Provenance*, arXiv (June 20, 2020), <https://arxiv.org/abs/2001.07886>.

¹³ See, e.g., Paul England, et al., *AMP: Authentication of Media via Provenance*, arXiv.

¹⁴ See *Creating the Standard for Digital Content Attribution*, Content Authenticity Initiative, <https://contentauthenticity.org/>; and *Project Origin: Protecting Trusted Media*, Project Origin, <https://www.originproject.info/about>.

¹⁵ Robert Williams has described how policy makers face an “innovation-security conundrum,” one aspect of which is “the worry that data privacy and national security are increasingly interconnected. Data (and data networks) can be exploited in ways that threaten security, but they also form the lifeblood of technological innovation on which both economic growth and national security depend.” Robert D. Williams, *Crafting a Multilateral Technology and Cybersecurity Policy*, Brookings at 1 (Nov. 2020), <https://www.brookings.edu/wp-content/uploads/2020/11/Robert-D-Williams.pdf>.

¹⁶ Ellen Nakashima, *With a Series of Major Hacks, China Builds a Database on Americans*, Washington Post (June 5, 2015), https://www.washingtonpost.com/world/national-security/in-a-series-of-hacks-china-appears-to-building-a-database-on-americans/2015/06/05/d2af51fa-0ba3-11e5-95fd-d580f1c5d44e_story.html.

¹⁷ Another example of an adversary acquiring significant data on U.S. individuals is the hack of the credit reporting agency Equifax. Press Release, Department of Justice, *Chinese Military Personnel Charged with Computer Fraud, Economic Espionage and Wire Fraud for Hacking into Credit Reporting Agency Equifax* (Feb. 10, 2020), <https://www.justice.gov/opa/pr/chinese-military-personnel-charged-computer-fraud-economic-espionage-and-wire-fraud-hacking>; Aruna Viswanatha, et al., *Four Members of China's Military Indicted Over Massive Equifax Breach*, Wall Street Journal (Feb. 11, 2020), <https://www.wsj.com/articles/four-members-of-china-s-military-indicted-for-massive-equifax-breach-11581346824>.

¹⁸ See, e.g., Drew Harwell & Eva Dou, *Huawei Tested AI Software That Could Recognize Uighur Minorities and Alert Police, Report Says*, Washington Post (Dec. 8, 2020), <https://www.washingtonpost.com/technology/2020/12/08/huawei-tested-ai-software-that-could-recognize-uighur-minorities-alert-police-report-says/>; Hugh Harsono, *China's Surveillance Technology Is Keeping Tabs on Populations Around the World*, The Diplomat (June 18, 2020), <https://thediplomat.com/2020/06/chinas-surveillance-technology-is-keeping-tabs-on-populations-around-the-world/>.

¹⁹ See James Baker, *Counterintelligence Implications of Artificial Intelligence—Part III*, Lawfare (Oct. 10, 2018), <https://www.lawfareblog.com/counterintelligence-implications-artificial-intelligence-part-iii>.

²⁰ Yuan Yang & James Fontanella-Khan, *Grindr Sold by Chinese Owner After US National Security Concerns*, Financial Times (March 7, 2020), <https://www.ft.com/content/a32a740a-5fb3-11ea-8033-fa40a0d65a98>.

²¹ On privacy attacks, see Maria Rigaki & Sebastian Garcia, *A Survey of Privacy Attacks in Machine Learning*, arXiv (July 15, 2020), <https://arxiv.org/abs/2007.07646>.

²² The Committee on Foreign Investment in the United States (CFIUS) has the authority to review transactions that include “sensitive personal data of United States citizens that may be exploited in a manner that threatens national security.” For background, see Laura Jehl, *Spotlight on Sensitive Personal Data As Foreign Investment Rules Take Force*, The National Law Review (Feb. 18, 2020), <https://www.natlawreview.com/article/spotlight-sensitive-personal-data-foreign-investment-rules-take-force>. The National Counterintelligence and Security Center (NCSC) includes “sensitive government data, and personally-identifiable information” in its conception of key supply chain risks. See *Supply Chain Risk Management: Reducing Threats to Key U.S. Supply Chains*, NCSC at 3 (2020), <https://www.dni.gov/files/NCSC/documents/supplychain/20200925-NCSC-Supply-Chain-Risk-Management-tri-fold.pdf>.

²³ See, e.g., Graham Webster, *App Bans Won't Make U.S. Security Risks Disappear*, MIT Technology Review (Sept. 21, 2020), <https://www.technologyreview.com/2020/09/21/1008620/wechat-tiktok-ban-china-us-security-policy-opinion/>.

²⁴ Nicholas Duran, et al., *2018 Webroot Threat Report*, Webroot at 6 (2018), https://www-cdn.webroot.com/9315/2354/6488/2018-Webroot-Threat-Report_US-ONLINE.pdf.

²⁵ Gary J. Saavedra, et al., *A Review of Machine Learning Applications in Fuzzing*, arXiv (Oct. 9, 2019), <https://arxiv.org/pdf/1906.11133.pdf>; Isao Takaesu, *Machine Learning Security: DeepExploit*, GitHub (Aug. 29, 2019), https://github.com/13o-bbr-bbq/machine_learning_security/tree/master/DeepExploit; Marc Ph. Stoecklin, et al., *DeepLocker: How AI Can Power a Stealthy New Breed of Malware*, Security Intelligence (Aug. 8, 2018), <https://securityintelligence.com/deeplocker-how-ai-can-power-a-stealthy-new-breed-of-malware/>.

²⁶ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine (2019), <https://doi.org/10.17226/25488>; Nektaria Kaloudi & Jingyue Li, *The AI-Based Cyber Threat Landscape*, ACM Computing Surveys at 1-34 (Feb. 2020), <https://dl.acm.org/doi/abs/10.1145/3372823>; Ben Buchanan, et al., *Automating Cyber Attacks*, Center for Security and Emerging Technology (Nov. 2020), <https://cset.georgetown.edu/research/automating-cyber-attacks/>; Dakota Cary & Daniel Cebul, *Destructive Cyber Operations and Machine Learning*, Center for Security and Emerging Technology at 5-23 (Nov. 2020), <https://cset.georgetown.edu/research/destructive-cyber-operations-and-machine-learning/>.

²⁷ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine (2019), <https://doi.org/10.17226/25488>.

²⁸ The recent SolarWinds attack demonstrates deep vulnerabilities in our software supply chains. See *Joint Statement by the Federal Bureau of Investigation (FBI), the Cybersecurity and Infrastructure Security Agency (CISA), and the Office of the Director of National Intelligence (ODNI)*, Office of the Director of National Intelligence (Dec. 16, 2020), <https://www.dni.gov/index.php/newsroom/press-releases/item/2175-joint-statement-by-the-federal-bureau-of-investigation-fbi-the-cybersecurity-and-infrastructure-security-agency-cisa-and-the-office-of-the-director-of-national-intelligence-odni>.

²⁹ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine (2019), <https://doi.org/10.17226/25488>; Ben Buchanan, et al., *Automating Cyber Attacks*, Center for Security and Emerging Technology at 3 (Nov. 2020), <https://cset.georgetown.edu/research/automating-cyber-attacks/>.

³⁰ *Cyberspace Solarium Commission Report*, U.S. Cyberspace Solarium Commission (March 2020), <https://www.solarium.gov/report>.

³¹ See recommendation 5.4 in *Cyberspace Solarium Commission Report*, U.S. Cyberspace Solarium Commission at 87 (March 2020), <https://www.solarium.gov/report>.

³² *Adversarial AI Threat Matrix: Case Studies*, GitHub (last accessed Jan. 10, 2021), <https://github.com/mitre/advm/threatmatrix/blob/master/pages/case-studies-page.md>. For more on applications of adversarial AI, see Naveed Akhtar & Ajmal Mian, *Threat of Adversarial Attacks on Deep Learning in Computer Vision: A Survey*, IEEE (March 28, 2018), <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8294186>.

³³ Adversarial AI is about what can be done to AI systems. The science of protecting and defending AI applications against attacks is called “AI Assurance.” The science of attacking each technological component of AI is called “Counter-AI.”

³⁴ Ram Shankar Siva Kumar & Ann Johnson, *Cyberattacks Against Machine Learning Systems Are More Common Than You Think*, Microsoft Security (Oct. 22, 2020), <https://www.microsoft.com/security/blog/2020/10/22/cyberattacks-against-machine-learning-systems-are-more-common-than-you-think/>.

³⁵ It has been estimated that less than 1% of AI R&D funding is directed toward the security of AI systems. See Nathan Strout, *The Three Major Security Threats to AI*, Center for Security and Emerging Technology (Sept. 10, 2019), <https://cset.georgetown.edu/article/the-three-major-security-threats-to-ai/>.

³⁶ Ram Shankar Siva Kumar, et al., *Adversarial Machine Learning—Industry Perspectives*, arXiv at 2 (May 21, 2020), <https://arxiv.org/pdf/2002.05646.pdf>.

Chapter 1 - Endnotes

³⁷ There are various ongoing public and private efforts including, for instance, the MITRE-Microsoft adversarial ML framework. See Ram Shankar Siva Kumar & Ann Johnson, *Cyberattacks Against Machine Learning Systems Are More Common Than You Think*, Microsoft Security (Oct. 22, 2020), <https://www.microsoft.com/security/blog/2020/10/22/cyberattacks-against-machine-learning-systems-are-more-common-than-you-think/>; *Adversarial AI Threat Matrix: Case Studies*, MITRE (last accessed Jan. 10, 2021), <https://github.com/mitre/advmthreatmatrix/blob/master/pages/case-studies-page.md>.

³⁸ For a similar recommendation, see Michèle Flournoy, et al., *Building Trust Through Testing*, WestExec Advisors at 27 (Oct. 2020), <https://cset.georgetown.edu/wp-content/uploads/Building-Trust-Through-Testing.pdf>. (Flournoy, et al., argue for “a national AI and ML red team as a central hub to test against adversarial attacks, pulling together DoD operators and analysts, AI researchers, T&E [Central Intelligence Agency (CIA), Defense Intelligence Agency (DIA), National Security Agency (NSA)], and other IC components, as appropriate. This would be an independent red-teaming organization that would have both the technical and intelligence expertise to mimic realistic adversary attacks in a simulated operational environment.”)

³⁹ BGI built and operates China National GeneBank, the Chinese government’s national genetic database. It also is a major global supplier of COVID-19 testing, which potentially provides access to large international genetic data sets; by June 30, 2020, it had supplied more than 35 million test kits to 180 countries, including the United States, and built 58 testing labs in 18 countries. See Kirsty Needham, Special Report: *COVID Opens New Doors for China’s Gene Giant*, Reuters (Aug. 5, 2020), <https://www.reuters.com/article/us-health-coronavirus-bgi-specialreport/special-report-covid-opens-new-doors-for-chinas-gene-giant-idUSKCN2511CE>.

⁴⁰ John Wertheim, *China's Push to Control Americans' Health Care Future*, 60 Minutes (Jan. 31, 2021), <https://www.cbsnews.com/news/china-us-biodata-60-minutes-2021-01-28/?ftag=CNM-00-10aab7d&linkId=110169507>; Kirsty Needham, *Special Report: COVID Opens New Doors for China's Gene Giant*, Reuters (Aug. 5, 2020), <https://www.reuters.com/article/us-health-coronavirus-bgi-specialreport/special-report-covid-opens-new-doors-for-chinas-gene-giant-idUSKCN2511CE>.

⁴¹ See Richard Pérez-Peña, *What Is Novichok, the Russian Nerve Agent Tied to Navalny Poisoning?*, New York Times (Sept. 2, 2020), <https://www.nytimes.com/2020/09/02/world/europe/novichok-skripal.html>; *2020 Adherence to and Compliance with Arms Control, Nonproliferation, and Disarmament Agreements and Commitments (Compliance Report)*, U.S. Department of State at Pt. V (2020), https://2017-2021.state.gov/2020-adherence-to-and-compliance-with-arms-control-nonproliferation-and-disarmament-agreements-and-commitments-compliance-report-2//index.html#_Toc43298166.

⁴² See Chapter 16 of this report for additional recommendations pertaining to the nexus of AI and biotechnology.

Chapter 2: Foundations of Future Defense

AI-Enabled Future Defense



**Build the
Technical
Backbone**



**Accelerate
Adoption of
Existing Digital
Technologies**



**Train and Educate
Warfighters**




**Invest in Next-
Generation
Capabilities**



**Democratize AI
Development**


The U.S. military has enjoyed military-technical superiority over all potential adversaries since the end of the Cold War. Now, its technical prowess is being challenged, especially by China and Russia. Senior military leaders have warned that if current trend lines are not altered, the U.S. military will lose its military-technical superiority in the coming years.¹ Artificial intelligence (AI) is a key aspect of this challenge, as both of our great power competitors believe they will be able to offset our military advantage using AI-enabled systems and AI-enabled autonomy. In the coming decades, the United States will win against technically sophisticated adversaries only if it accelerates adoption of AI-enabled sensors and systems for command and control, weapons, and logistics.

The Department of Defense (DoD) must set an ambitious goal. By 2025, the foundations for widespread integration of AI across DoD must be in place. Those foundations include a common digital infrastructure that is accessible to internal AI development teams and critical industry partners alike, a digitally literate workforce, and modern AI-enabled business practices that improve efficiency. All are prerequisites to achieving a state of military AI readiness, which is discussed in Chapter 3 of this report.



“By 2025, the foundations for widespread integration of AI across DoD must be in place.”

DoD lags far behind the commercial sector in integrating new and disruptive technologies such as AI into its operations. Pockets of excellence started to emerge in 2017 when Project Maven was launched with the aim to simplify work for intelligence analysts by recognizing objects in video footage captured by drones and other platforms.² Other promising initiatives are occurring in defense labs and agencies, and proof-of-concept demonstrations are ongoing in service-level tests.³ However, visionary technologists and warfighters largely remain stymied by antiquated technology, cumbersome processes, and incentive structures that are designed for outdated or competing aims.⁴ Successes are usually based on workarounds—in spite of the system.



“... visionary technologists and warfighters largely remain stymied by antiquated technology, cumbersome processes, and incentive structures that are designed for outdated or competing aims.”

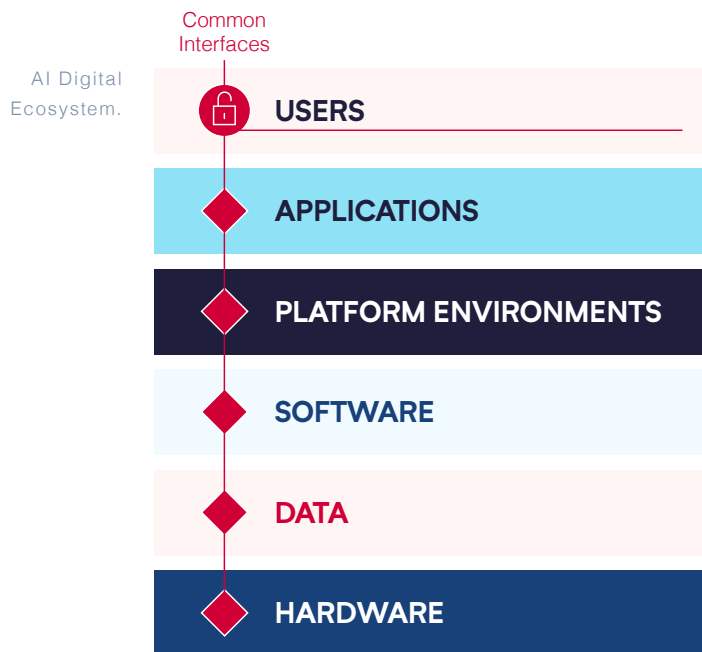
The obstacles to integrating AI are many. DoD has long been hardware-oriented toward ships, planes, and tanks. It is now trying to make the leap to a software-intensive enterprise. Spending remains concentrated on legacy systems designed for the industrial age and Cold War.⁵ Many Departmental processes still rely too much on PowerPoint and manually driven work streams. The data that is needed to fuel machine learning (ML) is currently stovepiped, messy, or often discarded. Platforms are disconnected. Acquisition, development, and fielding practices largely follow rigid, sequential processes, inhibiting early and continuous experimentation and testing critical for AI. Even promising AI programs have not yet delivered as hoped and often remain bound to proprietary software and data storage of commercial vendors. Steps such as building the cloud infrastructure necessary to scale AI applications proceed slowly. Data-sharing agreements and software updates that take hours or days in industry turn into months-long delays. Service members at every level lack the technical education and experience to employ AI.

Meanwhile, bureaucracy hinders partnerships with technology firms and critical efforts to expand the National Security Innovation Base.⁶ The prospect of bureaucratic snarls deters companies from working with DoD; it is economically irrational for many startups to even try. Traditional defense companies will continue to play a central role in building and integrating large systems for AI-enabled warfare.⁷ However, even these contractors, who have the resources and expertise to navigate the system, face process and technical roadblocks that slow efforts to build and integrate AI systems.

As a result, change will not be easy. It will require a Secretary of Defense who focuses the Department on speeding the adoption of new technologies, and a dedicated Steering Committee on Emerging Technology to drive implementation and align priorities between the DoD and the Intelligence Community. The Secretary should direct action in five areas:

Recommendation

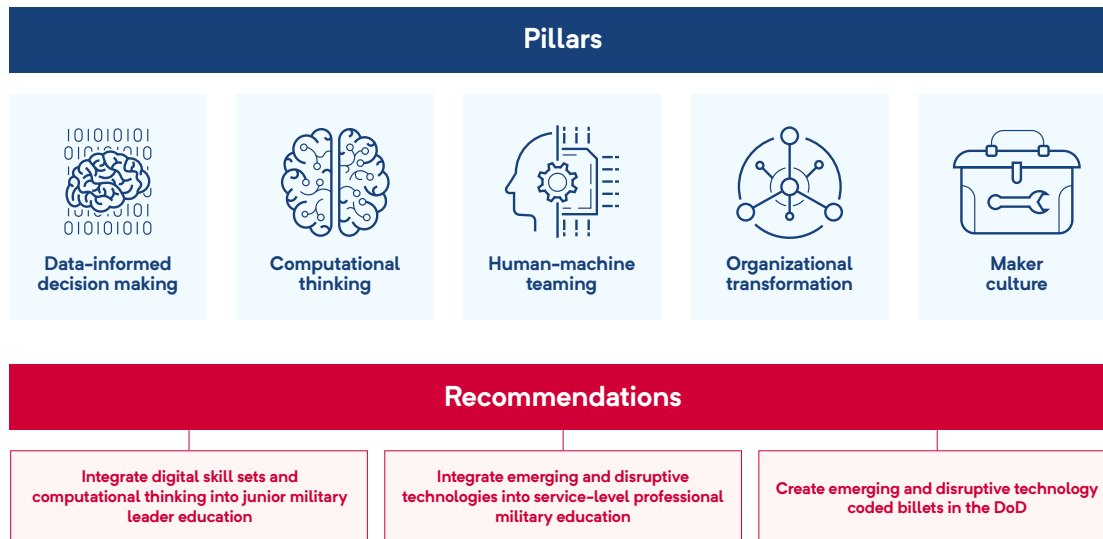
1. Build the technical backbone. DoD should make foundational investments to support a Department-wide technical infrastructure for ubiquitous development and fielding of AI. It took a promising first step in 2020 with the issuance of a DoD Data Strategy.⁸ However, the Department lacks the modern digital ecosystem, collaborative tools and environments, and broad on-demand access to shared AI resources that it needs to integrate AI across the organization.⁹ The Department should avoid reinventing core infrastructure for each new AI-driven program or capability, and it should look to leverage and interoperate with proven solutions from the Intelligence Community (IC) wherever possible. A broader platform that could be used across the Department would enable more dynamic development and employment of AI and would more efficiently utilize scarce technical expertise.¹⁰



The Secretary of Defense should direct the establishment of a DoD-wide digital ecosystem. The Secretary should require that all new joint and service programs adhere to the design of this ecosystem and that, wherever possible, existing programs become interoperable with it by 2025.¹¹ Key elements should include:

- Data architecture composed of a secure, federated system of distributed repositories linked by a data catalog and appropriate access controls¹² that facilitates finding, accessing, and moving desired data across the DoD.¹³
- Packaged AI environments¹⁴ that enable agile and iterative AI capabilities development,¹⁵ testing, fielding, and updating in support of a diverse set of stakeholders.¹⁶
- A marketplace of shared AI resources¹⁷ that builds upon federated repositories of data, software, and trained models,¹⁸ along with pre-negotiated computing and storage services from a pool of vetted cloud providers.
- A bolstered network and communications backbone to provide bandwidth to support transport and data fusion, secure processing, continuous development and fielding of AI applications, and software system integration at all levels.
- Common interfaces that allow swift integration of mission-oriented investments.

“Warfighters cannot change the way they fight without also changing the way they think.”



Train and Educate Warfighters.

2. Train and educate warfighters. Warfighters cannot change the way they fight without also changing the way they think. Most service members only use the powerful computers they have to create PowerPoint presentations, build spreadsheets, or send emails. Our service members need to develop core competencies in building, using, and responsibly teaming with machine systems to recognize AI's potential for building a faster and more effective force. In particular, they need to know:

Recommendation

- How to use data in decision-making in ways that complement intuition and experience.
- How to use information processing agents and how to get a computer to perform calculations and analytics that could not be done efficiently by a human.
- How to develop and thrive in a “maker” culture that encourages continuous contact and regular experimentation with and development of new tools.
- How to move toward a “teammate model” for interacting with autonomous systems and navigate issues of delegated authority, observability, predictability, directability, and trust.
- How to bring organizations into the AI era—including when and how to integrate AI-related tasks into priority missions, allocate resources to build and maintain the AI stack, oversee new systems, and support the careers of technical experts.

To improve training and education along these lines, DoD should:

- Identify service members who excel at computational thinking during the accession process;
- Invest in upskilling its workforce through self-guided education courses and coding language incentives;
- Teach junior leaders about problem curation, the AI lifecycle, data collection and management, probabilistic reasoning and data visualization, and data-informed decision-making as part of their pre-commissioning requirements and initial training;
- Integrate emerging and disruptive technology training into professional military education courses; and
- Create emerging technology coded billets and an emerging technology certification program comparable to the joint billet and qualification system.

Recommendation

3. *Accelerate the adoption of existing digital technologies.* DoD has largely relied on workarounds to adopt new technologies, while the core acquisition processes remain sclerotic. There are some bright spots, including the release of the Department's tailorable acquisition framework, contracting resources,¹⁹ and approaches taken by certain programs within the Air Force.²⁰ The Department must scale these innovative practices and take further steps to align acquisition workforce training, program incentives, budget, and organizational structures to better support the delivery of digitally enabled capabilities.

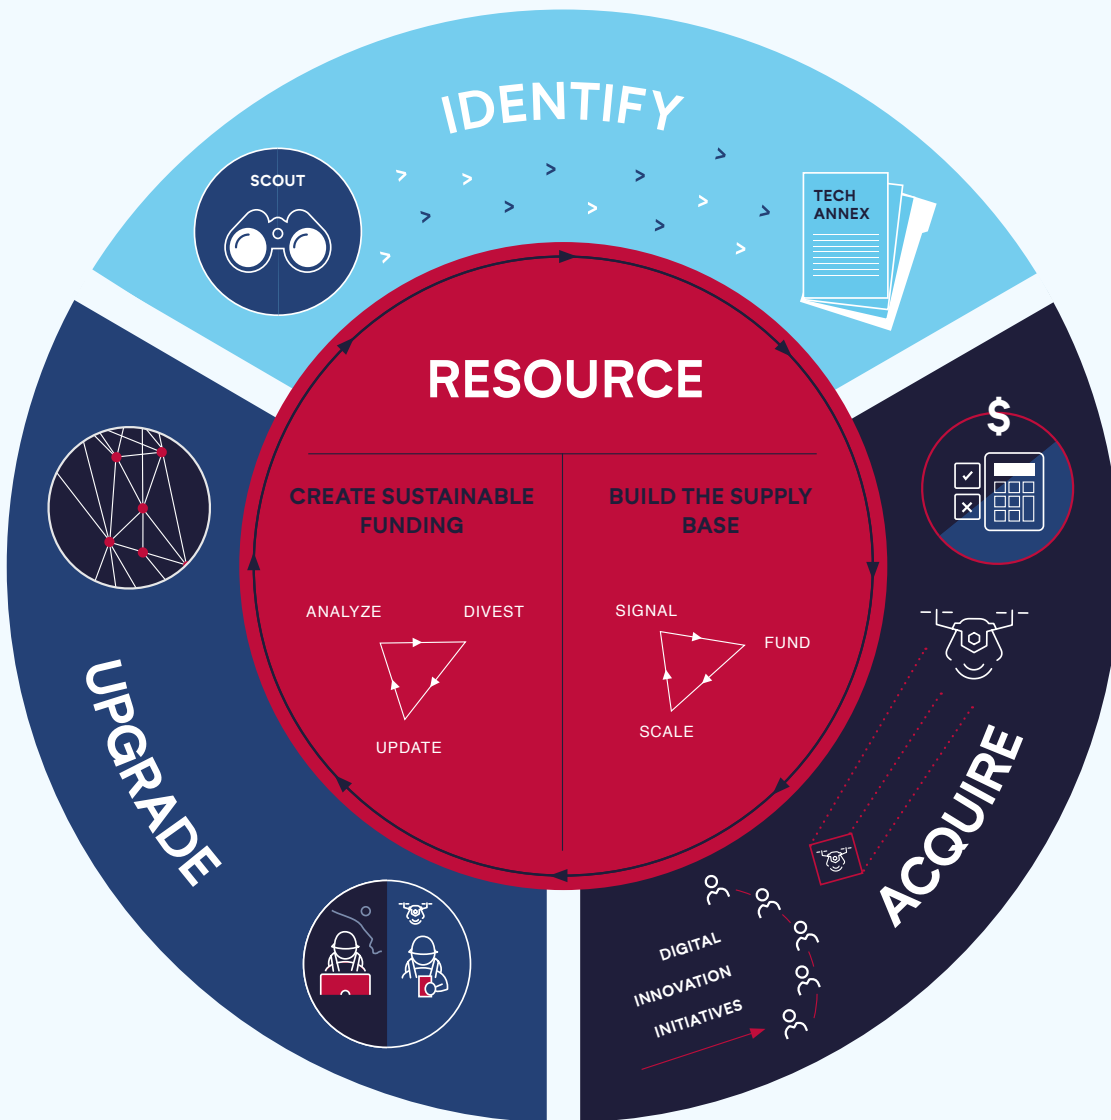
A number of the Department's digital innovation initiatives have delivered results,²¹ but they are uncoordinated and under-resourced. DoD signaling of technology priorities is ad hoc and is not supported by a track record of significant DoD investments in digital technology with non-traditional vendors. As a result, national security AI applications attract less private-market investment. The Department should focus on four actions:

- **Integrate commercial AI to optimize core business processes.** DoD should embrace proven commercial AI applications and incentivize their use to generate labor and cost savings, speed administrative actions, and inform decision-making.²² As a critical first step, DoD should prioritize construction of enterprise data sets across core administration areas.
- **Network digital innovation initiatives to scale impact.** Pockets of bottom-up innovation need to be married with top-down leadership. The Department should harmonize its innovation initiatives to carry out a coordinated go-to-market strategy for commercial technology solutions. The Under Secretary of Defense for Research and Engineering, working closely with the Under Secretary of Defense for Acquisition and Sustainment, the military services and other headquarters counterparts, should provide strategic direction for this effort.
- **Expand use of specialized acquisition pathways and contracting approaches.** DoD should accelerate efforts to train acquisition professionals on the full range of available options for acquisition and contracting and incentivize their use for AI and digital technologies.²³

- Update the budget and oversight processes.** DoD's resource allocation process is nearly identical to what was put in place in 1961. It is incompatible with AI and other digital technologies. DoD and Congress should institute reforms that enable the advancement of software and digital technologies by accounting for speed, uncertainty, experimentation, and continuous upgrades.

An integrated and strategic approach to technology that aligns the process, incentives, and organizational culture of the DoD and the National Security Innovation Base as a pipeline to resource, prioritize, acquire and iterate capabilities critical to sustain the competitive advantage

Delivering AI at Speed and Scale.



“At every level, technologists, operators, and domain experts should function as integrated teams.”

Recommendation

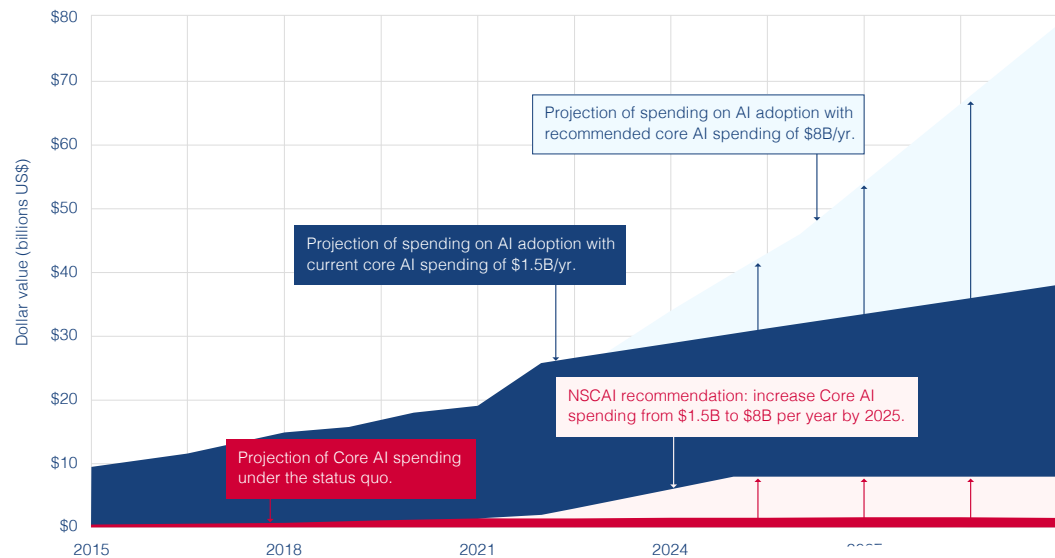
4. *Democratize AI development.* The Department must promote bottom-up AI development.²⁴ At every level, technologists, operators, and domain experts should function as integrated teams.²⁵ This would facilitate user feedback and improve trust and confidence in AI systems. DoD should:

- **Designate the Joint Artificial Intelligence Center (JAIC) as the Department’s AI Accelerator.** The JAIC cannot identify every potential use for AI in the Department, but it can and should serve as a central hub of AI expertise. In this “accelerator” model, JAIC would coordinate with relevant acquisition, technology, and governance offices to inform strategy; develop AI applications that address shared challenges at the Combatant Commands; and provide resources that enable distributed AI development across the Department and the military services.²⁶

Enhanced AI R&D Investment, FY 2015-2030

Source: Govini and NSCAI

Enhanced AI R&D Investment, FY 2015-2030.



This figure illustrates the correlation between R&D investment in Core AI technologies and AI adoption projected to the year 2030. Two scenarios are represented in this figure. In the first, the DoD maintains its current level of investment in core AI (~1.5B/year). In the second scenario, the DoD increases its

investments in core AI to \$8B/year. A significant increase in core AI spending is required to drive the rate of AI adoption higher.

NSCAI staff teamed with two external partners to analyze historical and planned DoD investments in AI RDT&E. The source data for the analysis is DoD's annual RDT&E budget expenditures (for FY2015 – FY 2020) and annual RDT&E budget requests (for FY2021-FY2025). For the methodology employed and lessons learned from this work, see Analysis of DoD RDT&E Investments in AI, NSCAI (on file with the Commission). Disclaimer: We believe this analysis yields important insights into general trends in AI spending and solutions for better future analyses, but caution that quality issues in the source data detailed in our on file report mean that the spending level estimates presented contain significant, difficult to estimate margins of error.

AI-enabled programs develop (in the case of RDT&E programs) and field (in the case of procurement programs) the gamut of DoD warfighting and business systems, incorporating Core AI applications for analyzing, automating, communicating, maneuvering, monitoring, sensing, and many other tasks. While AI spending is usually a small percentage of these programs, their system's performance may be critically dependent upon the incorporation of core AI.


AI-enabling programs include technologies such as cloud computing and advanced microelectronics required to support the deployment of effective AI capabilities at scale.

- **Establish integrated AI delivery teams at each Combatant Command.** These commands have specific operational needs that routinely outpace centralized development. AI delivery teams should be embedded at each Combatant Command and capable of supporting the full lifecycle of AI development and fielding, including data science, engineering, testing, and production—leveraging common resources through the digital ecosystem.²⁷ Teams should include forward-deployable components to act as the local interface with operational units.²⁸

5. *Invest in next-generation capabilities.* DoD leaders anticipate flat or declining defense budgets for the coming years.²⁹ Despite potential budgetary pressures, DoD must continue accelerating its modernization programs by prioritizing emerging and disruptive technologies such as AI.³⁰

Recommendation

- **Fund AI research and development.** The Department should commit to spending at least 3.4% of its budget on science and technology and allocate at least \$8 billion toward AI R&D annually.³¹ Additional resources should be focused on organizations with significant AI expertise, such as the Defense Advanced Research Projects Agency (DARPA), the Office of Naval Research (ONR), the Air Force Office of Scientific Research (AFOSR), the Army Research Office (ARO), and the service laboratories.



“To make AI ubiquitous throughout its business processes and military systems, DoD must make tough budget tradeoffs and prioritize modernization.”

- **Retire legacy systems ill-equipped to compete in AI-enabled warfare.** To make AI ubiquitous throughout its business processes and military systems, DoD must make tough budget tradeoffs and prioritize modernization.³² DoD should pursue a balanced approach to update existing systems with leading-edge technologies to buy time for investments in longer-term bets. Further, to guard against bias in favor of defending the status quo, DoD should require an evaluation of AI alternatives prior to funding Major Defense Acquisition Programs (MDAP).³³
- **Produce a technology annex to the National Defense Strategy.** To link DoD’s technology investment strategy to future operational needs, the annex should include roadmaps for designing, developing, fielding, and sustaining critical technologies that are needed to address the operational challenges identified in the strategy.

Chapter 2 - Endnotes

¹ General Joseph Dunford, then Chairman of the Joint Chiefs of Staff, testified in 2017 that “The U.S. military’s competitive advantage against potential adversaries is eroding [...] I assess that within five years we will lose our ability to project power; the basis of how we defend the homeland, advance U.S. interests, and meet our alliance commitments.” Posture Statement of General Joseph Dunford, Chairman of the Joint Chiefs of Staff before the Senate Armed Services Committee, *Senate Armed Services Budget Hearing* at 2 (June 13, 2017), https://www.armed-services.senate.gov/imo/media/doc/Dunford_06-13-17.pdf.

² *Big Data at War: Special Operations Forces, Project Maven, and Twenty-First-Century Warfare*, Modern War Institute (Aug. 25, 2020), <https://mwi.usma.edu/big-data-at-war-special-operations-forces-project-maven-and-twenty-first-century-warfare/>; Cheryl Pellerin, *Project Maven to Deploy Computer Algorithms to War Zone by Year’s End*, DoD (July 21, 2017), <https://www.defense.gov/Explore/News/Article/Article/1254719/project-maven-to-deploy-computer-algorithms-to-war-zone-by-years-end/>. Project Maven now includes detecting, classifying, and tracking objects within full motion video images (e.g., person, vehicle, and weapon) and other AI algorithms for text-based projects. *PE 0305245D8Z: Intelligence Capabilities and Innovation*, Office of the Secretary of Defense (Feb. 2019), https://www.dacis.com/budget/budget_pdf/FY20/RDTE/D/0305245D8Z_187.pdf.

³ For example, the Army’s Project Convergence exercise in September 2020 demonstrated use of AI at multiple stages of the targeting process. Jen Judson & Nathan Strout, *At Project Convergence, the US Army Experienced Success and Failure—and It’s Happy About Both*, Defense News (Oct. 12, 2020), <https://www.defensenews.com/digital-show-dailies/ausa/2020/10/12/at-project-convergence-the-us-army-experienced-success-and-failure-and-its-happy-about-both/>. The Air Force has held similar exercises, most notably as part of its efforts associated with the Advanced Battle Management System—the technical infrastructure which will support the DoD’s Joint All-Domain Command and Control concept. Theresa Hitchens, *ABMS Demo Proves AI Chops For C2*, Breaking Defense (Sept. 3, 2020), <https://breakingdefense.com/2020/09/abms-demo-proves-ai-chops-for-c2/>.

⁴ This includes the traditional process by which concepts of operation interact with technology development. Chapter 3 of this report offers recommendations to adapt this approach and ensure that technological advancements inform concepts as much as concepts drive technology development.

⁵ As one observer has noted: “While DoD’s investment accounts have grown substantially in the last three years, this growth has been highly concentrated in buying systems from existing production lines and doing prototypes of military systems.” Testimony of Andrew Hunter, Director, Defense-Industrial Initiatives Group, CSIS, before the U.S. House of Representatives Armed Services Committee, *Hearing on DoD’s Role in Competing with China* at 6 (Jan. 15, 2020), https://armedservices.house.gov/_cache/files/5/8/5818cc1f-b86f-4dca-8aee-10ca788e6f43/9F4A03ABF1DEAB747AF2D1302087A426.20200115-hasc-andrew-hunter-statement-vfinal.pdf.

⁶ The National Defense Strategy highlights the importance of the National Security Innovation Base in maintaining the Department’s technological advantage. *Summary of the 2018 National Defense Strategy*, U.S. Department of Defense at 3 (2018), <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>. The Center for Strategic and International Studies offers a useful definition of the term, noting that the “[National Security Innovation Base] is a significant expansion in scope [...] compared to the traditional concept of the defense industrial base” and includes tech firms out of innovation hubs such as Silicon Valley, Boston, and Austin. See Andrew Hunter, *A Strategic Approach to Defense Investment*, CSIS (March 26, 2018), <https://www.csis.org/analysis/strategic-approach-defense-investment>.

⁷ “The largest six prime defense suppliers (Lockheed Martin, Boeing, Northrop Grumman, Raytheon, General Dynamics, and BAE Systems) [...] represented 32 percent of all DoD prime obligations in 2019.” *Fiscal Year 2020: Industrial Capabilities*, U.S. Department of Defense at 40 (Dec. 23, 2020), https://www.businessdefense.gov/Portals/51/USA002573-20%20ICR_2020_Web.pdf?ver=o3D76uGwxcg0n0Yxvd5k-Q%3d%3d.

⁸ The strategy lays the foundation for the Department to treat data as a strategic asset and details the goals to make DoD data visible, accessible, understandable, linked, trustworthy, interoperable, and secure. *Executive Summary: DoD Data Strategy*, U.S. Department of Defense (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

Chapter 2 - Endnotes

⁹ In recent years the Department has made promising initial steps to establish managed services constructs for platforms, cloud infrastructure, and software development. For example, the Air Force's CloudOne and PlatformOne offerings (<https://software.af.mil/dsop/services/>); the Navy's Black Pearl (<https://blackpearl.us/>); and the Army's Coding Repository and Transformation Environment (CReATE). Further, the Office of the Secretary of Defense has built a data management platform, ADVANA, with the goal to establish it as the single authoritative source for audit and business data analytics. See Written Statement for the Record of David L. Norquist, Deputy Secretary of Defense before the U.S. Senate Armed Services Committee Subcommittee on Readiness at 6 (Nov. 20, 2019), https://www.armed-services.senate.gov/imo/media/doc/Norquist_11-20-19.pdf.

¹⁰ Components of this platform are underway as a result of the Joint Artificial Intelligence Center (JAIC)'s Joint Common Foundation initiative—particularly the marketplace of shared AI resources including data, algorithms, and trained AI models.

¹¹ Use of a common technical infrastructure will vastly improve DoD's ability to ensure interoperability and increase the effectiveness of the joint force. However, it is important to note that even without such critical technical infrastructure, the Department is taking important policy steps to drive interoperability and AI readiness for programs designed to meet joint capability needs. See Aaron Mehta, *Hyten to Issue New Joint Requirements on Handling Data*, Defense News (Sept. 23, 2020), <https://www.defensenews.com/pentagon/2020/09/23/hyten-to-issue-new-joint-requirements-on-handling-data/>. Chapter 3 of this report outlines additional recommendations for achieving a state of military AI readiness by 2025.

¹² Secured access to data sets as well as other shared resources should be managed by user- and role-based authentication facilitated by an end-to-end identity, credential, and access management infrastructure.

¹³ This hinges on implementation of the DoD's new data strategy. *Executive Summary: DoD Data Strategy*, U.S. Department of Defense (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

¹⁴ These are platform environments with ready-made workflows that can be tailored and launched depending on user type (e.g., researcher, industry partner, operator) and use case (e.g., development, TEVV [test, evaluation, validation, and verification], fielding).

¹⁵ In other words, the DevSecOps application lifecycle. "DevSecOps improves the lead time and frequency of delivery outcomes through enhanced engineering practices, promoting a more cohesive collaboration between Development, Security, and Operations teams as they work towards continuous integration and delivery." *Understanding the Differences Between Agile & DevSecOps—From a Business Perspective*, GSA (last accessed Jan. 1, 2021), <https://tech.gsa.gov/guides/understanding-differences-agile-devsecops/>.

¹⁶ Stakeholders could include embedded development teams working at the tactical edge; private-sector partners contributing pre-trained models; academic researchers working on open, relevant challenge problems; government science and technology (S&T) researchers working within a service lab; or international partners co-developing interoperable AI capabilities.

¹⁷ Shared AI resources should be managed with continuous Authorization to Operate (ATO) frameworks and with mandated default ATO reciprocity across the Department.

¹⁸ Similar to or relying upon the platform delivery and features of Git (<https://git-scm.com>), GitHub (<https://github.com>), and GitLab (<https://about.gitlab.com>).

¹⁹ The Pentagon acquisition office's Adaptive Acquisition Framework and Contracting Cone mark important steps by the Department to promote the use of alternate authorities for acquisitions and contracting. These include, for example, other transaction authorities, middle-tier acquisitions, rapid prototyping and rapid fielding, and specialized pathways for software acquisition.

²⁰ For example, the Air Force's Advanced Battle Management System (ABMS), which is managing systems intended to support the new Joint All-Domain Command and Control concept as a portfolio and based heavily on experimentation to drive innovation and an iterative approach to requirements. Notably, the Department of Defense Appropriations Bill for Fiscal Year 2021 expresses concern with various aspects of the Air Force's approach, including the "absence of firm requirements, acquisition

strategy, or cost estimate” and system of systems integration. See H. Rept. 116-453, at 294-295 (July 16, 2020), <https://www.congress.gov/116/crpt/hrpt453/CRPT-116hrpt453.pdf>.

²¹ The term “digital innovation initiatives” is used here to describe the various entities across the Office of the Secretary of Defense and the military services—such as the Defense Innovation Unit (DIU), AFWERX, NavalX, and Army Applications Laboratory (AAL)—that are focused on bridging the gap with the commercial technology sector, especially startups and non-traditional vendors, and accelerating the delivery of best-of-breed technology solutions.

²² The Defense Innovation Unit (DIU) is currently pursuing a number of AI projects to optimize business processes in the DoD, ranging from using AI-driven Robotic Process Automation to reduce labor costs for the Army Comptroller to improving Air Force readiness with AI-driven predictive maintenance and leveraging AI-constructed knowledge graphs to rapidly identify supply chain risks for the Defense Intelligence Agency. See *JAIC Partners with DIU on AI/ML Models to Resolve Complex Financial Errors*, JAIC (Oct. 1, 2020), https://www.ai.mil/blog/10_01_20-jaic_partners_with_diu_on_ai_ml_models_to_resolve_complex_financial_errors.html; *U.S. Defense Department Awards C3.ai \$95M Contract Vehicle to Improve Aircraft Readiness Using AI*, Business Wire (Jan. 15, 2020), <https://www.businesswire.com/news/home/20200115005413/en/US-Defense-Department-Awards-C3.ai-95M-Contract-Vehicle-to-Improve-Aircraft-Readiness-Using-AI>; *Accelerate.AI Accelerates Growth and Product Adoption with Defense Innovation Unit Contract*, Accrete.ai (April 23, 2020), <https://blog.accrete.ai/newsroom/accrete.ai-wins-million-dollar-contract-with-the-defense-innovation-unit>.

²³ As an example, DIU uses several acquisition pathways and contracting strategies that could help improve both the adoption and operational relevance of AI solutions and also expand the National Security Innovation Base. DIU pioneered the Commercial Solutions Opening with Army Contracting Command—New Jersey, which leverages section 2371b of title 10 U.S.C. Other Transaction authority to create a “fast, flexible, and collaborative” contract vehicle to prototype capabilities for the Department. DIU has also used Section 2374a of title 10 U.S.C. Prize Challenge authority to advance various AI-related priorities for DoD and the broader AI research community.

²⁴ The Department-wide digital infrastructure described above is critical to enabling this approach, but structural changes are also required to maximize its utility.

²⁵ There are notable examples of warfighter-technologist pairings within DoD, such as the Air Force’s software factories and the forward-deployed tactical data teams used by Special Operations and Army Futures Command. They found that partnering technologists (such as data scientists) with operators or analysts at the tactical edge: 1) significantly reduces the time it typically takes a contractor to understand the problem set and deploy a solution; 2) incentivizes iterative development techniques and fast-fielding of minimum viable products that yield higher-impact solutions on an accelerated timeline; and 3) generates increased buy-in to data and AI technologies as critical mission enablers. NSCAI Engagements (Nov. 2020). To ensure U.S. forces maintain overmatch in the long-term, DoD must scale this user-centered development.

²⁶ Important offices for coordination with the JAIC include but are not limited to USD(R&E), USD Acquisition & Sustainment (USD(A&S)), Director Operational Test & Evaluation (DOT&E), and the DoD Chief Information Officer (CIO) and Chief Data Officer (CDO). Within USD(R&E), DIU is a key enabler of the JAIC that pursues a project-based approach by transitioning commercial prototypes for specific applications. The JAIC currently serves the Combatant Commands through its Component Mission Initiatives (CMIs), including a Mission Initiative for Joint Warfighting Operations. See *Mission Initiatives*, JAIC (last accessed Dec. 28, 2020), https://www.ai.mil/mi_joint_warfighting_operations.html.

²⁷ Such applications could be developed by other Combatant Commands, Service software factories, or the JAIC and discoverable via the recommended digital ecosystem. Each Combatant Command should ensure that the AI delivery teams are staffed with the appropriate talent to manage the full lifecycle of AI solutions, including in disciplines such as data science, AI testing and model training, software engineering, product management, and full stack development.

²⁸ As an example, both Army Futures Command (AFC) and Army Special Operations Command (USASOC) use a model known as “tactical data teams.” This model brings AI/ML expertise forward to the field in the form of three- to six-person teams to build AI solutions for real-time operational problems. Executed by a small business, Striveworks, under contract with AFC and USASOC, they are currently supporting efforts in Central Command and Indo-Pacific Command Areas of Responsibility.

Chapter 2 - Endnotes

²⁹ Jim Garamone, *Chairman Discusses Future Defense Budgets*, U.S. Department of Defense (Dec. 3, 2020), <https://www.defense.gov/Explore/News/Article/Article/2433856/chairman-discusses-future-defense-budgets/>.

³⁰ *Summary of the 2018 National Defense Strategy*, U.S. Department of Defense at 6 (2018), <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>.

³¹ The Defense Science Board has proposed the level of 3.4% in the past to mirror typical practices in the private sector. *Department of Defense Research, Development, Test, and Evaluation (RDT&E): Appropriations Structure*, Congressional Research Service at 12 (Oct. 7, 2020), <https://fas.org/sgp/crs/natsec/R44711.pdf>.

³² The Future of Defense Task Force report similarly stated that “policy makers, industry, and the Pentagon must work together to identify trade-offs within the defense apparatus to include legacy systems and operations, which will allow for investment in technology and operational concepts to address future challenges.” *Future of Defense Task Force Report 2020*, House Armed Services Committee at 18 (Sept. 23, 2020), <https://armedservices.house.gov/cache/files/2/6/26129500-d208-47ba-a9f7-25a8f82828b0/424EB2008281A3C79BA8C7EA71890AE9.future-of-defense-task-force-report.pdf>.

³³ This should utilize wargaming, experimentation, and live-virtual-constructive environments wherever feasible, and should mandate interoperability with the digital ecosystem. This point echoes the Future of Defense Task Force, which recommended that every Major Defense Acquisition Program (MDAP) should be required “to evaluate at least one AI or autonomous alternative prior to funding.” *Id.* at 7.

Chapter 3: AI and Warfare

AI Ready by 2025



Top-Down Leadership



Innovative Concepts



Advanced Technologies and R&D



AI-Readiness Performance Goals



AI-Enabled Allies and Partners

Even with the right artificial intelligence (AI)-ready technology foundations in place, the U.S. military will still be at a battlefield disadvantage if it fails to adopt the right concepts and operations to integrate AI technologies. Throughout history, the best adopters and integrators, rather than the best technologists, have reaped the military rewards of new technology.¹ The Department of Defense (DoD) should not be a witness to the AI revolution in military affairs, but should deliver it with leadership from the top, new operating concepts, relentless experimentation, and a system that rewards agility and risk.

A new warfighting paradigm is emerging because of AI. Our competitors are making substantial investments to take advantage of it. This idea has been called “algorithmic” or “mosaic” warfare²; China’s theorists have called it “intelligentized” war.³ All of these terms capture, in various ways, how a new era of conflict will be dominated by AI and pit algorithms against algorithms. Advantage will be determined by the amount and quality of a military’s data, the algorithms it develops, the AI-enabled networks it connects, the AI-enabled weapons it fields, and the AI-enabled operating concepts it embraces to create new ways of war.

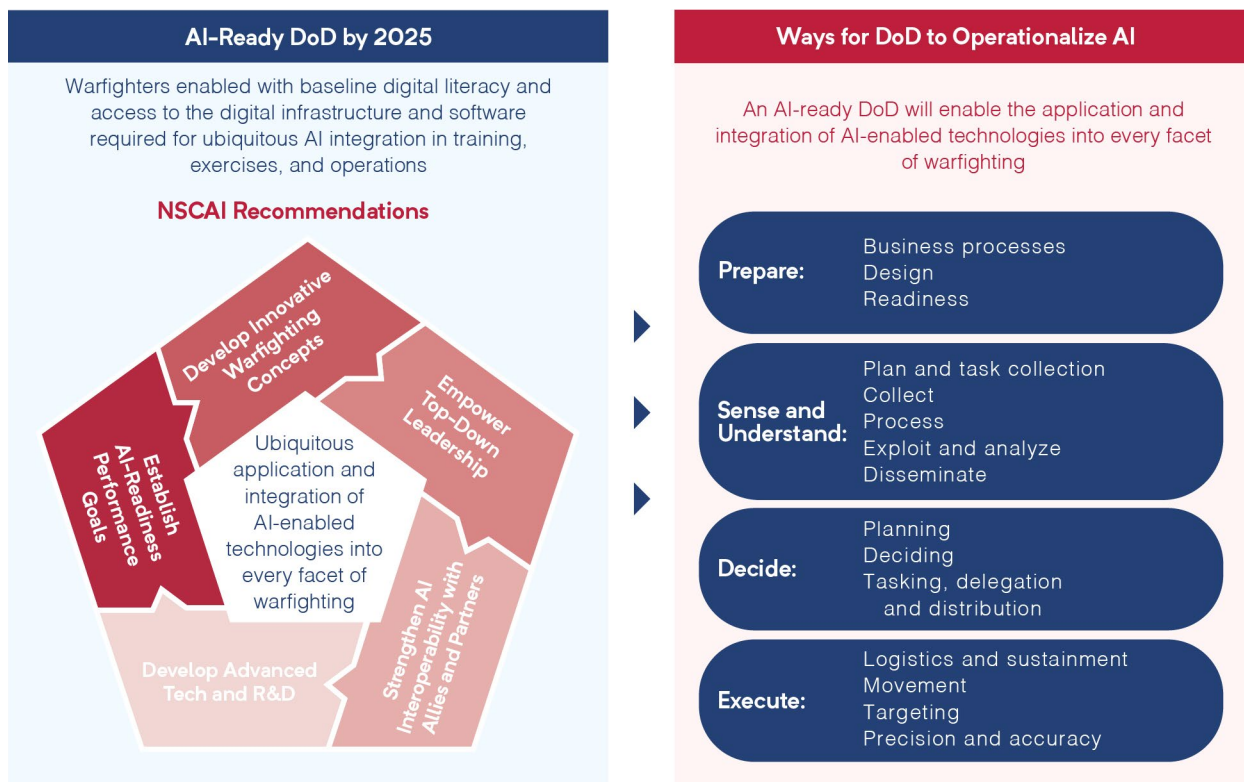
Today’s DoD is trying to execute an AI pivot, but without urgency. Despite pockets of imaginative reform and a few farsighted leaders, DoD remains locked in an Industrial Age mentality in which great-power conflict is seen as a contest of massed forces and monolithic platforms and systems. The emerging ubiquity of AI in the commercial realm and the speed of digital transformation punctuate the risk of not pivoting fast enough. The Department must act now to integrate AI into critical functions, existing systems, exercises, and wargames to become an AI-ready force by 2025. Simultaneously, DoD must develop more creative warfighting concepts that are paired with investments in future AI-enabled technologies to continuously out-innovate potential adversaries. If our forces are not equipped with AI-enabled systems guided by new concepts that exceed those of their adversaries, they will be outmatched and paralyzed by the complexity of battle.

An AI-Ready DoD by 2025:

Warfighters enabled with baseline digital literacy and access to the digital infrastructure and software required for ubiquitous AI integration in training, exercises, and operations.

“The Department must act now to integrate AI into critical functions, existing systems, exercises and wargames to become an AI-ready force by 2025.”


AI and Warfare



To compete, deter, and, if necessary, fight and win in future conflicts requires wholesale adjustments to operational concepts, technologies, organizational structures, and how we integrate allies and partners into operations. It will also require risk-based assessments of both the benefits and drawbacks of widespread integration of AI-enabled capabilities, to include future autonomous weapon systems. Lastly, it will require a willingness to engage in bilateral and multilateral dialogues with our allies and partners to urge them to make similar AI pivots to ensure future interoperability.

How AI Will Change Warfare.


AI-enabled warfare will not hinge on a single new weapon, technology, or operational concept; rather, it will center on the application and integration of AI-enabled technologies into every facet of warfighting. AI will transform the way war is conducted in every domain from undersea to outer space, as well as in cyberspace and along the electromagnetic spectrum. It will impact strategic decision-making, operational concepts and planning, tactical maneuvers in the field, and back-office support. In this new kind of warfare, traditional confines of the battlefield will be expanded through AI-enabled micro-targeting, disinformation, and cyber operations, as described in Chapter 1 of this report. AI will reshape many attributes of war, such as its speed, tempo, and scale; the relationships service members have with machines; the persistence with which the battlefield can be monitored; and the discrimination and precision with which targets can be attacked. There will be a premium on speed and accuracy in developing knowledge, acting, and reacting as the conflict unfolds.



“AI-enabled warfare will not hinge on a single new weapon, technology, or operational concept; rather, it will center on the application and integration of AI-enabled technologies into every facet of warfighting.”

DoD AI RDT&E Investments by Warfighting Domains, FY 2015–2025

Source: Govini



DoD AI RDT&E Investments.

DoD's AI investments are well distributed across the various warfighting domains (land, naval, air, space, cyber, electromagnetic spectrum, information), with over 25% of investments in multi-domain applications of AI, signaling AI's potential for integrating multi-domain operations. Investments in AI applications for Space operations more than quadrupled from FY2019 to FY2025, from \$500M to \$2.2B, increasing from 3% to almost 9% of AI-enabled investment.

Note the spending levels presented in figure represent estimates based on an analysis of DoD RDT&E budget documents for FY2021-FY2025. See Analysis of DoD RDT&E Investments in AI, NSCAI (on final with the Commission). Due to inherent quality issues in the source data, estimates presented contain significant, difficult to estimate margins of error.

AI will make the process of finding and hitting targets of military value faster and more efficient. It will also increase accuracy of target identification and minimize collateral damage. Currently, this process generally involves passing data in a serial fashion from a sensor, through a series of humans, to a platform that can shoot at the target. AI will help automate some of the intermediate stages of the decision process. AI will also create opportunities for more advanced processes that would operate more akin to a web, fusing multiple sensors and platforms to manage complex data flows and transmitting actionable information to human operators and machines across all domains.⁴

In war, many of the military uses of AI will complement, rather than supplant, the role of humans. AI tools will improve the way service members perceive, understand, decide, adapt, and act in the course of their missions. However, new concepts for military operations will also need to account for the changing ways in which humans will be able to delegate increasingly complex tasks to AI-enabled systems. In the near term, this will be managed through the military's principle of "mission command," which stresses decentralized execution and disciplined initiative by subordinates who follow a commander's intent. This human-centric approach to fighting should remain the standard for the foreseeable future. But as AI continues to advance into the cognitive and neuromorphic domain, and human-machine teaming becomes more sophisticated, the military will need to develop more imaginative concepts and organizational constructs that take full advantage of AI technologies without relinquishing the principles that undergird mission command.

Prepare

Business processes. Robotic Process Automation and AI-enabled analysis can generate significant savings, speed administrative actions, and provide decision-makers with superior insights into core business processes such as finance, budget, contracting, travel, and human resources.

Design. AI will support a holistic system-of-systems approach to developmental force design via digital engineering, digital twins, and modeling and simulation to enable a more comprehensive understanding of system vulnerabilities and adjacent capabilities, concepts, and technologies.

Readiness. AI will enhance training by relieving the cognitive burden of doing repetitive tasks that can be performed better by machines. AI will be prevalent in all exercises and wargames and will enhance the military's ability to train in live, virtual, and constructive environments.

Plan and task collection. Through automation, AI-enabled systems will optimize tasking and collection for platforms, sensors, and assets in near-real time in response to dynamic intelligence requirements or changes in the environment.

Sense and Understand

Collect. At the tactical edge, "smart" sensors will be capable of pre-processing raw intelligence and prioritizing what data to transmit and store, which will be especially helpful in degraded or low-bandwidth environments.

Process. AI-enabled natural language processing, computer vision, and audiovisual analysis can vastly reduce manual data processing. AI can also be used to automate data conversion such as translations and decryptions, accelerating the ability to derive actionable insights.

Exploit and analyze. AI-enabled tools have the potential to augment filtering, flagging, and triage across multiple data sets. Such tools can identify connections and correlations more efficiently and at a greater scale than human analysts and can flag those findings and the most important content for human analysis. AI will improve indications and warnings for military leaders.

- AI can fuse data from multiple sources, types of intelligence, and classification levels to produce accurate predictive analysis in a way that is not currently possible.
- Advances in speech-to-text transcription and language analytics now enable reading comprehension, question answering, and automated summarization of large quantities of text.

Disseminate. AI will be able to automatically generate machine-readable versions of intelligence products and disseminate them at machine speed so that computer systems across the IC and the military can ingest and use them in real time without manual intervention.

Decide

Planning. AI decision-support applications will utilize modeling and simulation algorithms and real-time data sets to optimize planning options.

Deciding. AI will integrate command-and-control networks and compress the speed of finding and attacking targets of military value.

Tasking, delegation, and distribution. Edge processing enhanced by delegated authorities will allow frontline units to operate in a coordinated manner with minimal to no communications. AI techniques like machine learning, and rule-based models will support network resiliency.

Execute


Logistics and sustainment. AI-enabled predictive analytics, optimization, and tracking will improve efficiency and effectiveness across all facets of logistics. Intelligent systems will aid in the development of courses of action for routine and contingency logistics and sustainment operations. Robotic process automation will streamline human-centric maintenance and supply chain workflows.

Movement. AI will enhance the ability of commanders to maneuver, position, and protect units and forces. AI will help network and coordinate movements of autonomous swarms via human-machine and machine-machine teaming.

Targeting. AI-enabled systems will expand a single targeting chain into a complex targeting web that considers numerous variables across units and domains.

Precision and accuracy. Through AI-enabled smart weapons and autonomous platforms, AI will enable the military to be more precise and discern friendly forces, non-combatants, and adversary targets with greater accuracy.

This list of how AI might transform warfighting principles and capabilities—as well as others like it—is by no means exhaustive. Innovation will lead to future capabilities that are unknowable at present and will only become clearer in time.



“Innovation will lead to future capabilities that are unknowable at present and will only become clearer in time.”

Stronger Together.

If the United States wants to fight with AI, it will need allies and partners with AI-enabled militaries and intelligence agencies. Uneven adoption of AI will threaten interoperability and the political cohesion and resiliency of U.S. alliances.⁵ As it deepens and expands conventional defense arrangements across the globe—especially in Europe and the Indo-Pacific—the United States should incorporate AI and emerging technology into coordinated defense and intelligence activities. Given the dual-use nature of many software-based capabilities, DoD will need more flexibility to work with civilian agencies, companies, and research institutions in partner nations.


Promote AI interoperability and the adoption of critical emerging technologies among allies and partners, including the Five Eyes, the North Atlantic Treaty Organization (NATO), and across the Indo-Pacific. This should include:

- Enhancing existing Five Eyes AI-related defense and intelligence efforts.
- Supporting NATO efforts to accelerate agreements on architectures and standards, develop allied technical expertise, and pursue coalition AI use cases for exercises and wargames.
- Fostering the Joint Artificial Intelligence Center (JAIC)'s International AI Partnership for Defense as a critical vehicle to further AI defense and security cooperation.⁶
- Creating an Atlantic-Pacific Security Technology Partnership to improve military and intelligence capabilities and interoperability across European and Indo-Pacific allies and partners.

Recommendation

“Uneven adoption of AI will threaten military interoperability, and the political cohesion and resiliency of U.S. alliances.”

AI-enabled
Alliances and
Partnerships.



Achieve a State of Military AI Readiness by 2025.

To reach this goal, the DoD should:

Recommendation

Drive organizational reform through top-down leadership. Senior civilian and military officials should set clear priorities and direction, empower subordinates, and accept higher uncertainty and risk in pursuing new technologies. Specifically, DoD should:

- Establish a high-level Steering Committee on Emerging Technology, tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence⁷;
- Ensure that the JAIC Director remains a three-star general or flag officer with significant operational experience who reports directly to the Secretary of Defense or Deputy Secretary of Defense;
- Appoint the Under Secretary of Defense for Research and Engineering as the co-chair and chief science advisor to the Joint Requirements Oversight Council; and

- Assign an AI Operational Advocate on the staff of every Combatant Command. This officer would perform a similar role to that played by the Staff Judge Advocate. He or she would be an expert in AI systems, advise the commander and staff on the capabilities and limitations of AI systems, and identify when AI-enabled systems are being used inappropriately.

*Develop innovative operational concepts that integrate new warfighting capabilities with emerging technologies.*⁸ These concepts should strive for seamless interoperability across the military services and across operational domains. The concept developers should work closely with technologists to articulate how the military could fight most effectively in future scenarios, and they should assume that AI-enabled capabilities will be ubiquitous on future battlefields. These concepts can also drive future investments.

Recommendation

*By the end of 2021, establish AI and digital readiness performance goals.*⁹ To achieve more substantial integration of AI across DoD, the Deputy Secretary of Defense should:

Recommendation









- Direct DoD components to assess military AI and digital readiness through existing readiness management forums and processes. The Tri-Chaired Steering Committee on Emerging Technology should work closely with the Under Secretary of Defense for Personnel and Readiness and the Joint Staff to ensure the identified AI readiness criteria are incorporated into the military services' readiness reporting and resourcing strategies.
- Direct the military services to accelerate review of specific skill gaps in AI to inform recruitment and talent-management strategies.¹⁰
- Direct the military services—in coordination with the Under Secretary of Defense for Acquisition and Sustainment, the Joint Staff, the Defense Logistics Agency, and the JAIC—to prioritize integration of AI into logistics and sustainment systems wherever possible.
- Integrate AI into major wargames and exercises to promote field-to-learn approaches to technology adoption. Operators need persistent interaction with AI-enabled capabilities early in the development cycle to generate critical feedback on how they function and how they impact the mission. Widespread experimentation will advance both concept development and the performance of the technology.¹¹
- Incentivize experimentation with AI-enabled applications through the Warfighting Lab Incentive Fund, which could be overseen by the proposed Tri-Chaired Steering Committee.¹²

Define a joint warfighting network architecture by the end of 2021. The key objective of this joint warfighting network should be a secure, open-standards systems network that supports the integration of AI applications at operational levels and across domains.¹³ It should be accessible by all of the military services and encompass several elements, including command and control networks; data transport, storage, and secure processing; and weapon system integration. The technical infrastructure for the network should be supported by best practices in digital engineering.¹⁴ It should also be interoperable with the digital ecosystem described in Chapter 2 of this report.¹⁵

Recommendation

Invest in priority AI R&D areas that could support future military capabilities, including the following:

Recommendation

Research Area	Time Horizon		Key Challenges	Category
 Future of Teaming	Near: Understanding and designing support for human-machine interdependence. Proper management of interdependence to achieve smooth and efficient coordination of activity.	Long: Reasoning over changing context to flexibly adapt teaming strategy for the best available team performance. Achieving team resilience through flexible adaptation.	Understanding interdependencies and the impact of dynamic adaptation on team performance.	Work with Humans
 Advanced Scene Understanding	Near: Ability to sense fundamental changes in the operating environment and alert the human operator while switching to a better-suited, environmentally-tuned perceptual model, if one exists	Long: Maintaining a perceptual model that supports actionable awareness and insight across a range of complex, dynamic environments and scenarios	Incorporating multi-source and multi-modal information from complex and changing environments	Sense & Perceive
 Intelligent Edge Devices, Computing & Networking	Near: Narrow AI applications within edge sensors, such as remote cameras positioned to monitor a highly contested space	Long: Autonomous edge devices that dynamically learn, share, and team with other devices, while exercising intelligent data collection, exploitation, and retention; mastery of domain-specific physical manipulation	Network limitation; size, weight, and power (SWaP)	Hardware, Devices, and Robotics
 Robust and Resilient AI	Near: Standard practice for exchanging trained AI models with tamper resistance and non-repudiation	Long: AI systems that are resilient on attack surfaces and able to learn securely via privacy-centric machine learning, including use of encryption	Many attack surfaces; addressing rise of adversarial machine-learning methods with robust learning; applying security techniques while maintaining high accuracy	Integrate & Assure
 Test & Evaluation, Verification & Validation (TEVV)	Near: Common framework for AI TEVV	Long: TEVV for fully autonomous AI systems that employ dynamic learning along with self-awareness and monitoring, and autonomous AI test ranges involving cohorts of humans and machines	Knowing how much and what types of testing are sufficient to determine an acceptable level of risk for a given use case	
 Integrated AI, Modeling & Simulation for Decision Support	Near: Decision support for highly constrained scenarios and environments	Long: Real-time decision support and course of action development for open-world environments with longer time horizons	Multi-modal data integration; assessing the predictive fidelity of simulation models	Learn & Reason
 Autonomous AI Systems	Near: Operate for relatively short periods of delegated autonomy from human operators in relatively unchanging and predictable environments while carrying out simple, independent tasks	Long: Longer periods of independent mission engagement with awareness and understanding of a dynamically changing operational environment, requiring continual assurance and self-monitoring, while carrying out complex mission sets involving multi-agent collaboration	Independent accomplishment of goals in environments that are complex, changing, and unpredictable; understanding how and when to engage with human operators	
 Toward More General Artificial Intelligence	Near: Growth in interpretability and explainability for narrow AI; methods for performing transfer learning; fine-tuning of models	Long: AI systems able to learn by engaging with the operational environment, make decisions based on contextual knowledge, and amass experiential knowledge	Unlocking multiple mysteries of human learning and reasoning; more general situational awareness and problem-solving	

Chapter 3 - Endnotes

¹ On military adoption, see, e.g., Michael C. Horowitz, *The Diffusion of Military Power: Causes and Consequences for International Politics*, Princeton University Press (2010).

² The Defense Advanced Research Projects Agency (DARPA) Mosaic warfare central concept is built around the “adaptability for U.S. forces and complexity or uncertainty for the enemy through the rapid composition and recombination of a more disaggregated U.S. military force using human command and machine control.” Bryan Clark, et al., *Mosaic Warfare: Exploiting Artificial Intelligence and Autonomous Systems to Implement Decision-Centric Operations*, CSBA at vi (Feb. 11, 2020), <https://csbaonline.org/research/publications/mosaic-warfare-exploiting-artificial-intelligence-and-autonomous-systems-to-implement-decision-centric-operations/publication/1>.

³ The People’s Liberation Army has developed a warfighting concept for what it calls “intelligentized operations” with AI at its core. Within this construct, China theorizes that in future conflict, the central contest will be between adversarial battle networks rather than traditional weapons platforms, and that information advantage and algorithmic superiority will be a determinant of victory. See Elsa Kania, *Chinese Military Innovation in Artificial Intelligence*, CNAS at 1 (June 7, 2019), <https://www.cnas.org/publications/congressional-testimony/chinese-military-innovation-in-artificial-intelligence> (testimony before the U.S.-China Economic and Security Review Commission).

⁴ See *Creating Cross-Domain Kill Webs in Real Time*, DARPA (Sept. 18, 2020), <https://www.darpa.mil/news-events/2020-09-18a>. See also *AI Fusion: Enabling Distributed Artificial Intelligence to Enhance Multi-Domain Operations & Real-Time Situational Awareness*, Carnegie Mellon University (2020), <http://www.cs.cmu.edu/~ai-fusion/overview>.

⁵ On military interoperability challenges related to AI, see Erik Lin-Greenberg, *Allies and Artificial Intelligence: Obstacles to Operations and Decision-Making*, Texas National Security Review (Spring 2020), <https://tnsr.org/2020/03/allies-and-artificial-intelligence-obstacles-to-operations-and-decision-making/>.

⁶ The AI Partnership for Defense, launched in September 2020, includes the United States and 12 partner nations: Australia, Canada, Denmark, Estonia, Finland, France, Israel, Japan, Norway, South Korea, Sweden, and the United Kingdom. It seeks to “provide values-based global leadership” on adoption of AI in the defense and security context and align “like-minded nations to promote the responsible use of AI, advance shared interests and best practices on AI ethics implementation, establish frameworks to facilitate cooperation, and coordinate strategic messaging on AI policy.” *Joint Statement*, AI Partnership for Defense (Sept. 15-16, 2020), https://www.ai.mil/docs/AI_PfD_Joint_Statement_09_16_20.pdf. The Partnership held its second formal dialogue in January 2021. *DoD Joint AI Center Facilitates Second International AI Dialogue for Defense*, JAIC (Jan. 27, 2021), https://www.ai.mil/news_01_27_21-dod_joint_ai_center_facilitates_second_international_ai_dialogue_for_defense.html.

⁷ The Commission acknowledges Section 236 of the Fiscal Year 2021 National Defense Authorization Act, which permits the Secretary of Defense to establish a steering committee on emerging technology and national security threats composed of the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, the Under Secretary of Defense for Intelligence and Security, the Under Secretary of Defense for Research and Engineering, the Under Secretary of Defense for Personnel and Readiness, the Under Secretary of Defense for Acquisition and Sustainment, the Chief Information Officer, and such other officials of the Department of Defense as the Secretary determines appropriate. However, the structure described in Sec. 236 does not include leadership from the Intelligence Community and will thus not drive the intended action. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁸ Notably, the National Defense Strategy emphasizes the need to “evolve innovative operational concepts” and “foster a culture of experimentation and calculated risk-taking.” Tighter coordination between concept writers and technologists would create a more dynamic cycle of technology development and integration. *Summary of the 2018 National Defense Strategy*, U.S. Department of Defense at 7 (2018), <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>.

Chapter 3 - Endnotes

⁹ “Readiness” is a key measure of military effectiveness and remains at the heart of budget, policy, and oversight debates on defense preparedness. In this context, DoD should establish key AI and digital readiness performance objectives to measure and drive Department and service accountability. See G. James Herrera, *The Fundamentals of Military Readiness*, Congressional Research Service at 2 (Oct. 2, 2020), <https://fas.org/sgp/crs/natsec/R46559.pdf>.

¹⁰ As noted in Chapter 6 of this report, there is already an identified need for the creation of digital corps, civilian and military AI and AI-related career fields, an expansion of recruiting pathways, and the creation of recruiting offices. The military services need to assess the number of personnel in those fields and structures, not the need to establish them.

¹¹ Although AI will be ubiquitous across all domains, the high-data volumes associated with the space, cyber, and information operations domains make use cases in those domains particularly well-suited for prioritized integration of AI-enabled applications in wargames, exercises, and experimentation.

¹² The Warfighting Lab Incentive Fund is intended to spur field experiments and demonstrations to “evaluate, analyze and provide insight into more effective ways of using current capabilities, and to identify new ways to incorporate technologies into future operations and organizations.” See Memorandum from the Deputy Secretary of Defense, *Warfighting Lab Incentive Fund and Governance Structure*, U.S. Department of Defense (May 6, 2016), https://defenseinnovationmarketplace.dtic.mil/wp-content/uploads/2018/02/DSD_memo.pdf.

¹³ The network envisioned is well-aligned with ongoing DoD efforts to embrace standards-driven interoperability, system adaptability, and data-sharing. See Memorandum from the Secretary of the Navy, Secretary of the Army, and Secretary of the Air Force for Service Acquisition Executives and Program Executive Officers, U.S. Department of Defense (Jan. 7, 2019), https://www.dsp.dla.mil/Portals/26/Documents/PolicyAndGuidance/Memo-Modular_Open_Systems_Approach.pdf.

¹⁴ Such as the goals and focus areas outlined in the DoD Digital Engineering Strategy; terms, any knowledge, and guidelines shared as part of the Digital Engineering Body of Knowledge; and incorporating Section 231 of the National Defense Authorization Act for Fiscal Year 2020, which requires the creation of a digital engineering capability to automate testing and evaluation. See *Department of Defense Digital Engineering Strategy*, Office of the Deputy Assistant Secretary of Defense for Systems Engineering (June 2018); see also Pub. L. 116-92, The National Defense Authorization Act for Fiscal Year 2020, 133 Stat. 1198 (2019). For a description of the Digital Engineering Body of Knowledge, see Andrew Monje, *Future Direction of Model-Based Engineering Across the Department of Defense*, U.S. Department of Defense (Jan. 27, 2020), <https://ac.cto.mil/wp-content/uploads/2020/05/RAMS-Monje-27Jan2020-Future.pdf>.

¹⁵ See the Chapter 2 Blueprint for Action for details on how this architecture should interact with the digital ecosystem.

Chapter 4: Autonomous Weapon Systems and Risks Associated with AI-Enabled Warfare

Mitigate Strategic Risks Associated with AI-Enabled Weapon Systems



**Continue
Rigorous TEVV
Procedures**



**Develop
International
Standards of
Practice**



**Discuss
Risks with
Competitors**



**Limit Specific
Applications**

World military powers both large and small are pursuing artificial intelligence (AI)-enabled and autonomous weapon systems. Such systems have the potential to help commanders make faster, better, and more relevant decisions. They will enable weapon systems to be capable of levels of performance, speed, and discrimination that exceed human capabilities. And they will enable hitherto impossible complex tasks. If properly designed, tested, and used, they could improve compliance with International Humanitarian Law (IHL)¹ by reducing the risk of accidental engagements, decreasing civilian casualties, minimizing collateral infrastructure damage, and allowing for detailed auditing of the decisions and actions of operators and their command chains. Although U.S. weapons platforms have utilized autonomous functionalities for more than eight decades,² AI technologies have the potential to enable novel, sophisticated offensive and defensive autonomous capabilities.

The increasing use of AI technologies in weapon systems has generated important questions regarding whether such systems are lawful, safe, and ethical. Those critical of using AI technologies in weapons argue that states should negotiate limits or restrictions on such systems and their use. There is also concern that autonomous weapon systems may make conflict escalation more likely, and debate continues over what steps are needed to ensure that such systems minimize the risk of unintended military engagements or inadvertent and uncontrollable conflict escalation. Since 2014, the United Nations Convention on Certain Conventional Weapons (CCW) has held meetings among states parties to discuss the technological, military, legal, and ethical dimensions of “emerging technologies in the area of lethal autonomous weapon systems (LAWS).”³ Specifically, it is examining whether autonomous technologies will be capable of complying with IHL and whether additional measures are necessary to ensure that humans maintain an appropriate degree of control over the use of force.

The Commission has consulted with civil society, academic organizations, and government agencies in studying the legal, ethical, and strategic questions that surround AI-enabled and autonomous weapon systems, including their potential military benefits and risks, possible ethical issues coming to the fore, international efforts to regulate them, and their compliance with IHL. The Commission offers the following four judgments to reflect its conclusions on these discussions.

Judgment 1: Provided their use is authorized by a human commander or operator, properly designed and tested AI-enabled and autonomous weapon systems have been and can continue to be used in ways which are consistent with IHL.


This judgment is grounded in several elements of IHL:

- **Distinction:** The principle of distinction holds that parties to an armed conflict must distinguish between civilians and combatants.⁴ Weapons with increasingly accurate AI-enabled target recognition systems have the potential to reduce cases of target misidentification, the leading cause of inadvertent engagements during combat operations, and thus reduce civilian casualties and collateral infrastructure damage.⁵
- **Proportionality:** The principle of proportionality prohibits attacks which would cause incidental loss of civilian life excessive to the anticipated military advantage.⁶ AI-enabled and autonomous weapon systems can and should also be designed to carry out operations in accordance with human judgments and directions regarding the proportionality of an attack. The moral reasoning involved in this calculus—weighing anticipated military advantage against potential civilian harm—remains the responsibility of a human commander.⁷
- **Accountability:** Ensuring accountability and command responsibility is essential to compliance with IHL. A human can and should be held accountable for the development, testing, use, and behavior of any autonomous weapon system, AI-enabled or otherwise. Autonomous weapon systems operate within the same general parameters as those used for human command and control systems, which are specifically designed to ensure accountability for actions and compliance with IHL. This is no different than for any other weapon system.⁸

NSCAI Judgments Regarding AI-Enabled and Autonomous Weapon Systems

- Provided their use is authorized by a human commander or operator, properly designed and tested AI-enabled and autonomous weapon systems have been and can continue to be used in ways which are consistent with IHL.
- Existing DoD procedures are capable of ensuring that the United States will field safe and reliable AI-enabled and autonomous weapon systems and use them in a manner that is consistent with IHL.
- There is little evidence that U.S. competitors have equivalent rigorous procedures to ensure their AI-enabled and autonomous weapon systems will be responsibly designed and lawfully used.
- The Commission does not support a global prohibition of AI-enabled and autonomous weapon systems.

The Commission endorses DoD's body of policy that states that human judgment must be involved in decisions to take human life in armed conflict. The kind of involvement necessary for humans to remain accountable for the use of autonomous weapon systems will vary depending on the time criticality of the situation as well as the operational context, circumstance, and type of weapon systems involved.⁹ It is incumbent upon states to establish processes which ensure that appropriate levels of human judgment are relied



“... human judgment must be involved in decisions to take human life in armed conflict.”

upon in the use of AI-enabled and autonomous weapon systems and that human operators of such systems remain accountable for the results of their employment.


Human accountability for the results of lethal engagements does not necessarily require human oversight of every step of an engagement process. Once a human authorizes an engagement against a target or group of targets, subsequent steps in the attack sequence can be completed autonomously without relinquishing human accountability. The exact number of steps in this sequence is dependent on the system’s technical capabilities and the context and must consider factors such as the uncertainty associated with the system’s behavior and potential outcomes, the magnitude of the threat, and the time available for action. For instance, an autonomous weapon system located in a rapidly changing environment, such as an urban setting, for an extended period, may require more frequent human authorization to ensure sufficient human accountability over autonomous actions than an equivalent system, operated for a similar amount of time, in a highly predictable and less populated environment—such as underwater or in space. This logic can and should be incorporated into the system’s design, testing, and operational planning. Taking these factors into consideration, when feasible and deemed necessary operation designs should include points of required human guidance amid a sequence of automated actions. At such points, a human must review the system’s status and authorize its next actions before the system’s mission can continue. A blanket decision to compel every discrete step in an engagement involving lethal force to be subject to explicit authorization by a human is neither realistic nor desirable. Indeed, such a policy could instead spur commanders to use less precise, unguided weapon systems that might result in greater levels of collateral damage.

Judgment 2: Existing DoD procedures are capable of ensuring that the United States will field safe and reliable AI-enabled and autonomous weapon systems and use them in a manner that is consistent with IHL.

DoD’s commitment to rigorous procedures for the development and use of autonomous weapon systems—as well as its commitment to strong AI ethical principles¹⁰—instills confidence that it will be able to field AI-enabled and autonomous weapon systems that are used lawfully. DoD has comprehensive processes for ensuring that the use of any weapon it fields is compliant with IHL and has a demonstrated commitment to operating within IHL, minimizing civilian casualties, and learning from its mistakes.¹¹ DoD has established

a cross-department legal group, the DoD Law of War Working Group, to “develop and coordinate law of war initiatives and issues, such as analysis regarding the legality of new means or methods of warfare under consideration by DoD components.”¹² This standing body is well positioned to examine implications for IHL as technology evolves over time. The International Committee on the Red Cross (ICRC) has lauded the strength and transparency of this system, listing the United States as one of eight countries that have “national mechanisms to review the legality of weapons and that have made the instruments setting up these mechanisms available to the ICRC.”¹³

In addition to baseline legal review, the Department has taken special precautions for autonomous weapon systems to ensure these systems undergo sufficient test and evaluation, verification and validation (TEVV). In 2012, DoD added to an extensive list of guiding directives and instructions regarding weapons development within the Department by publishing DoD Directive (DoDD) 3000.09, *Autonomy in Weapon systems*, which establishes DoD policy for the development and use of autonomous weapon systems. It requires that all systems be designed “to allow commanders and operators to exercise appropriate levels of human judgment over the use of force” and requires senior DoD leaders to approve any autonomous weapon with lethal capabilities first when development begins, and again before fielding.¹⁴ It also mandates any autonomous or semi-autonomous weapon that undergoes a revision to its operating state to undergo additional testing and evaluation. DoDD 3000.09 provides important definitions and baseline requirements for such systems and must be reviewed annually as technology evolves.¹⁵ Chapter 7 of this report provides specific recommendations on how the United States should adapt its TEVV policies and capabilities to ensure it retains justified confidence in AI-enabled systems.¹⁶



“The U.S. commitment to IHL is longstanding, and AI-enabled and autonomous weapon systems will not change this commitment.”


In addition, DoD’s command and control procedures to authorize target selection and employment of munitions are rigorous and designed to ensure compliance with IHL. Operational commanders in the field are directly supported by lawyers embedded at multiple levels to advise on decisions about the use of force. The U.S. commitment to IHL

is long-standing, and AI-enabled and autonomous weapon systems will not change this commitment.¹⁷ These same principles will be ingrained into the design of those weapons, demonstrated in TEVV, and maintained by commanders overseeing their deployment. DoD's policy for autonomy in weapon systems and its adoption of ethical principles for AI in 2020 further highlight and reinforce this commitment.¹⁸

Judgment 3: There is little evidence that U.S. competitors have equivalent rigorous procedures to ensure their AI-enabled and autonomous weapon systems will be responsibly designed and lawfully used.

Battlefield success may become increasingly dependent on AI performance, and AI-enabled weapons are likely to proliferate given the open-source and dual-use nature of AI. This could cause pressure to mount on states to rapidly field new and untested systems and algorithms. Such pressures could also tilt designs toward systems that react more quickly, limiting the amount of time available for effective human oversight on engagement decisions. U.S. competitors, particularly Russia and China, likely do not have equivalent operational and targeting procedures to ensure the use of such systems is compliant with IHL and to preserve human accountability over the use of lethal force. Russia and China also have not published anything equivalent to DoDD 3000.09, outlining their policies and processes governing the acquisition, development, testing, and deployment of autonomous weapon systems. Unlike in the United States, in Russia and China these processes are secret, if they exist at all.

U.S. competitors have demonstrated that they are unlikely to adhere to the same ethical and legal standards in developing and utilizing AI-enabled weapon systems. Russia in particular has historically demonstrated a willingness to deploy risky and under-tested weapon systems, and it has deployed poorly performing unmanned ground vehicles



“A global treaty prohibiting the development, deployment, or use of AI-enabled and autonomous weapon systems is not currently in the interest of U.S. or international security ...”

with limited autonomous functionalities in combat in Syria.¹⁹ China is not only actively pursuing increased autonomous functionality across a range of military systems, but it is also currently exporting armed drones with autonomous functionalities to other nations. This includes systems such as the Blowfish A3, which Ziyang, the system's manufacturer, advertises as capable of conducting autonomous, lethal, targeted strikes.²⁰

Judgment 4: The Commission does not support a global prohibition of AI-enabled and autonomous weapon systems.

A global treaty prohibiting the development, deployment, or use of AI-enabled and autonomous weapon systems is not currently in the interest of U.S. or international security and would be inadvisable to pursue for several reasons:

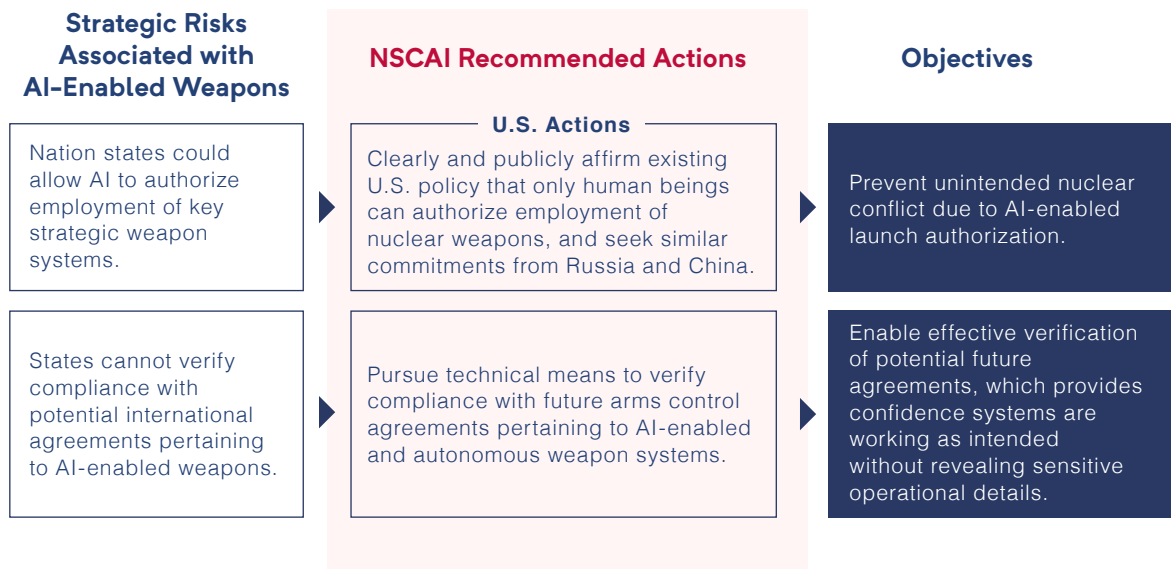
- First is the basic definitional problem. With respect to autonomous weapon systems, although the UN discussions about LAWS date back to 2014, states have yet to agree on a definition for them. This makes any treaty negotiation problematic, as it may be impossible to define the category of systems to be restricted in such a way that provides adequate clarity while not overly constraining existing U.S. military capabilities.
- Even if the definitional problem could be overcome, we judge that, at present, implementation of such an agreement would be impractical because compliance could not be verified. There is no feasible technical manner in which states could demonstrate to one another that specific weapon systems are or are not autonomous, or that they possess or lack certain capabilities. Doing so would require foreign inspectors to have short-notice access to the underlying code in weapon systems of concern. States are unlikely to agree to such an intrusive verification regime because revealing that information would create unacceptable risks to the security of their systems.
- Additionally, the effects of a prohibition agreement likely would run counter to U.S. strategic interests. Commitments from states such as Russia or China likely would be empty ones. Such an agreement would not serve the goal of putting political pressure on the states that are most likely to deploy autonomous weapon systems in unsafe and ethically concerning ways. Rather, the primary impact of an agreement would be to increase pressure on those countries that abide by international law, including the United States and its democratic allies and partners. Moreover, differing views on a prohibition among U.S. allies could deepen divisions among them on the employment of AI-enabled autonomous weapon systems. If U.S. allies joined an agreement while the United States did not, that divergence would likely hinder allied military interoperability.²¹

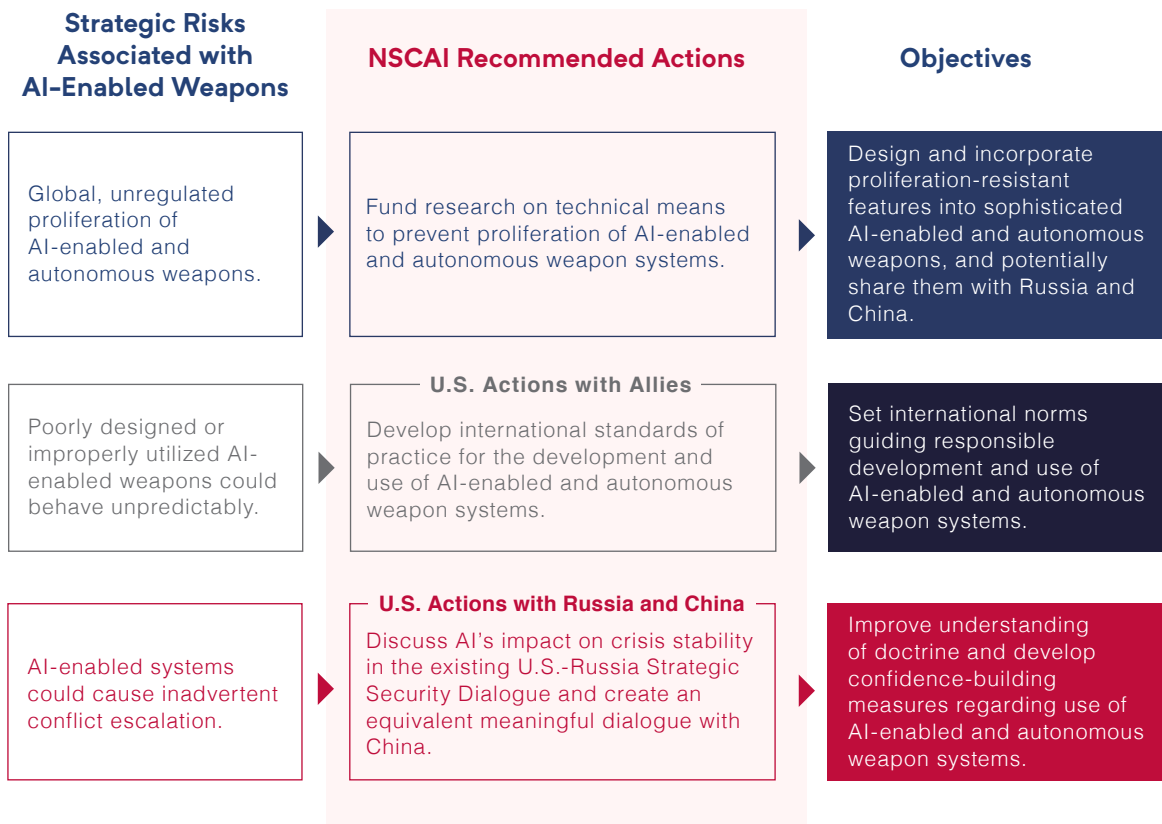
For these reasons, we believe the practical and strategic problems with a prohibition treaty outweigh potential benefits for the United States and its allies and partners, and therefore we support the current U.S. policy in opposition to such an agreement. However, this does not preclude other agreements or policies to address strategic risks associated with AI-enabled and autonomous weapon systems, or the future possibility of regulating specific types of technologies in AI-enabled and autonomous weapons technologies when such an agreement could be verifiable.

Recommendations to Mitigate Strategic Risks of AI.

While the Commission believes that properly designed, tested, and utilized AI-enabled and autonomous weapon systems will bring substantial military and even humanitarian benefit, the unchecked global use of such systems potentially risks unintended conflict escalation and crisis instability. The United States cannot assume that AI-enabled and autonomous weapon systems fielded by other countries will be developed, acquired, and fielded with the appropriate testing and verification to enable them to act as intended. Unintended escalations may occur for numerous reasons, including when systems fail to perform as intended, because of challenging and untested complexities of interaction between AI-enabled and autonomous weapon systems on the battlefield, and, more generally, as the result of machines or humans misperceiving signals or actions. AI-enabled systems will likely increase the pace and automation of warfare across the board, reducing the time and space available for de-escalatory measures. Beyond testing and robustness, we cannot assume that AI-enabled and autonomous weapons developed by other nations will be designed to behave in accordance with IHL.

Therefore, countries must take actions which focus on reducing risks associated with AI-enabled and autonomous weapon systems and encourage safety and compliance with IHL when discussing their development, deployment, and use. Such efforts should and must be led by the United States, which is uniquely situated to lead them given its technical expertise, military prowess, and clear and transparent policies and ethical principles governing the deployment and use of AI-enabled and autonomous weapon systems. The Commission presents the following five recommendations regarding actions the United States should take to mitigate risks associated with AI-enabled and autonomous weapon systems.





Clearly and publicly affirm existing U.S. policy that only human beings can authorize employment of nuclear weapons, and seek similar commitments from Russia and China. The United States should make a clear, public statement that decisions to authorize nuclear weapons employment must only be made by humans, not by an AI-enabled or autonomous system, and should include such an affirmation in the DoD's next Nuclear Posture Review.²² This would cement and highlight existing U.S. policy, which states that “[t]he decision to employ nuclear weapons requires the explicit authorization of the President of the United States.”²³ It would also demonstrate a practical U.S. commitment to employing AI and autonomous functions in a responsible manner, limiting irresponsible capabilities, and preventing AI systems from escalating conflicts in dangerous ways. It could also have a stabilizing effect, as it would reduce competitors’ fears of an AI-enabled, bolt-from-the-blue strike from the United States and could incentivize other countries to make equivalent pledges.

Recommendation

The United States should also actively press Russia and China, as well as other states that possess nuclear weapons, to issue similar statements. Although joint political commitments that only humans will authorize employment of nuclear weapons would not be verifiable, they could still be stabilizing, responding to a classic prisoner’s dilemma: as long as countries have confidence that others are not building risky command and control structures that have the potential to inadvertently trigger massive nuclear escalation, they would have less incentive to develop such systems themselves.²⁴ While this norm is widely accepted in the United States, it is unclear if Russia and China share the same strategic


“... countries must take actions which focus on reducing risks associated with AI-enabled and autonomous weapon systems, and encourage safety and compliance with IHL when discussing their development, deployment, and use. Such efforts should and must be led by the United States ...”

concerns. Public reports indicate that Russia previously installed a “dead hand” system to automate nuclear launch authorization,²⁵ and China’s representatives in Track II dialogues with the United States have been hesitant to state that China would make an equivalent commitment. If neither Russia nor China is willing to agree to such a proposal, the United States should mount a strong international pressure campaign to condemn this decision and highlight how Russia and China refuse to commit to responsible military uses of AI.

Recommendation

Discuss AI’s impact on crisis stability in the existing U.S.-Russia Strategic Security Dialogue (SSD) and create an equivalent meaningful dialogue with China. The Departments of State and Defense should discuss AI’s impact on crisis stability within the existing U.S.-Russia SSD and create an equivalent meaningful dialogue with China. The SSD is an interagency bilateral dialogue focused on reducing misunderstandings and misperceptions on key strategic issues and threats, as well as reducing the likelihood of inadvertent escalation. Although the dialogue has traditionally focused on nuclear arms control and doctrine, it has recently been used to also discuss emerging technologies and space security.²⁶ The United States has no equivalent dialogue with China, as China has resisted U.S. attempts to establish one for nearly a decade. However, within the last year there has been increasing evidence that China is interested in formal talks with the United States concerning AI-enabled military systems.²⁷ This interest should be cultivated and leveraged into establishing a U.S.-China SSD that includes the relevant military, diplomatic, and security officials from both sides.

Given that the United States, Russia, and China are all aggressively pursuing AI-enabled capabilities, and that Russia and China are likely to field AI-enabled systems that have undergone less rigorous TEVV than comparable U.S. systems and may be unsafe or unreliable, it is crucial to improve mutual understanding of each other's military doctrines, including with respect to AI-enabled and autonomous systems. The United States should use this channel to highlight how deploying unsafe systems could risk inadvertent conflict escalation, emphasize the need to conduct rigorous TEVV, and discuss where each side sees risks of a conventional conflict rapidly escalating in order to better anticipate future responses in a crisis.



“... it is crucial to improve mutual understanding of each other's military doctrines, including with respect to AI-enabled and autonomous systems.”

These dialogues could also plant the seeds for a future, standing dialogue exclusively focused on establishing practical and concrete confidence building measures surrounding AI-enabled and autonomous weapon systems. For instance, the United States, Russia, and China could work to develop an “international autonomous incidents agreement,” modeled after the 1972 Incidents at Sea Agreement, which would seek to define the “rules of the road” for behavior of autonomous military systems to create a more predictable operating environment and avoid accidents and miscalculations.²⁸ They could also agree to integrate “automated escalation tripwires” into systems that would prevent the automated escalation of conflict in specific scenarios without human intervention, to include nuclear weapons employment as noted above.

Work with allies to develop international standards of practice for the development, testing, and use of AI-enabled and autonomous weapon systems. The United States must work closely with its allies to develop standards of practice regarding how states should responsibly develop, test, and employ AI-enabled and autonomous weapon systems. This could build off of existing work, to include the 11 Guiding Principles agreed to by the LAWS Group of Governmental Experts (GGE) in 2019,²⁹ DoDD 3000.09, the DoD Ethical Principles for AI, and the NSCAI Key Considerations for Responsible Development and Fielding of AI.³⁰ As part of this effort, the DoD Law of War Working Group should meet regularly to review any future technical developments that pertain to autonomous weapon systems and IHL, and the tri-chaired Steering Committee on Emerging Technology (separately recommended

Recommendation

by the Commission in Chapter 3 of this report) should advise on how such future technical developments impact policy and national defense.

The outputs of both groups should inform future DoD engagements with both allies and competitors on AI-enabled and autonomous weapon systems. Obtaining allied consensus regarding standards for the development, testing, and use of such systems will set important norms regarding these systems, help to ensure they are developed and used safely, and further highlight the commitment of the United States and its allies to ethical and responsible uses of AI. The United States should also use these consultations to highlight the ways in which AI will become a crucial part of future military operations and develop common frameworks guiding the appropriate and responsible use of AI-enabled and autonomous weapon systems on the battlefield. This should seek to incentivize allies to invest in the digital modernization of their own forces while also highlighting the risks to military interoperability should any ally agree to join a treaty prohibiting LAWS.

Recommendation

Pursue technical means to verify compliance with future arms control agreements pertaining to AI-enabled weapon systems. The United States should actively pursue the development of technologies and strategies that could enable effective and secure verification of future arms control agreements involving uses of AI technologies. Although arms control of AI-enabled weapon systems is currently technically unverifiable, effective verification will likely be necessary to achieve future legally binding restrictions on AI capabilities. DoD and the Department of Energy (DoE) should spearhead efforts to design and implement technologies which could provide other countries confidence that an AI-enabled and autonomous weapon system is working as intended without revealing sensitive operational details. For instance, it could examine ways for AI-enabled weapons platforms to produce authenticatable records of operation, which could be spot-checked via international challenge inspections if noncompliant activity is suspected. Technical creativity will be necessary to enable any future international restrictions on AI capabilities without revealing sensitive information.

Recommendation

Fund research on technical means to prevent proliferation of AI-enabled and autonomous weapon systems. Controlling the proliferation of AI-enabled and autonomous weapon systems poses significant challenges given the open-source, dual-use, and inherently transmissible nature of AI algorithms.³¹ The proliferation of makeshift autonomous weapon systems which primarily utilize commercial components will be particularly difficult to control via regulation and will necessitate capable intelligence sharing and domestic law enforcement efforts to prevent their use by terrorists and other non-state actors. Regarding more sophisticated autonomous weapon systems, the United States should double down on efforts to design and incorporate proliferation-resistant features, such as standardized ways to prevent unauthorized users from utilizing such weapons, or reprogramming a system's functionality by changing key system parameters. DoD and DoE should fund technical research on such methods, and if appropriate, these methods could be shared with Russia and China, or potentially other countries, to prevent the proliferation or loss of control of certain AI-enabled autonomous weapon systems.³²

This report does not contain a separate Blueprint for Action for Chapter 4. This is because given the importance of the topic, the Commission chose to detail its arguments, recommendations, and the specific actions required to implement them directly in this chapter. Additionally, further detail on how the United States should adapt its TEVV policies to maintain confidence in AI systems can be found in Chapter 7 and its associated Blueprint for Action, and recommendations on relevant changes to DoD organizational structure can be found in Chapter 3.

Chapter 4 - Endnotes

¹ IHL is also referred to as the law of armed conflict (LOAC) and the law of war.

² Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, W.W. Norton & Co. at 39 (April 24, 2018).

³ *Background on Lethal Autonomous Weapon systems in the CCW*, United Nations (last accessed Jan. 11, 2021), [https://www.unog.ch/80256EE600585943/\(httpPages\)/8FA3C2562A60FF81C1257CE600393DF6?OpenDocument](https://www.unog.ch/80256EE600585943/(httpPages)/8FA3C2562A60FF81C1257CE600393DF6?OpenDocument).

⁴ *Distinction*, International Committee of the Red Cross (last accessed Jan. 15, 2021), <https://casebook.icrc.org/glossary/distinction>.

⁵ There is room for improvement in reducing target misidentification in U.S. military operations. In the Afghanistan war, for example, a study indicated that about half of all civilian casualty incidents caused by U.S. forces resulted from target misidentification. The use of AI-enabled systems to make more accurate targeting decisions is perhaps the principal way in which the proper employment of AI could make warfare more humane. Larry Lewis, *Redefining Human Control: Lessons from the Battlefield for Autonomous Control*, CNA at 4 (March 2018), https://www.cna.org/cna_files/pdf/DRM-2017-U-016281-Final.pdf.

⁶ *Proportionality*, International Committee of the Red Cross (last accessed Jan. 15, 2021), <https://casebook.icrc.org/glossary/proportionality>.

⁷ See Paul Scharre, *Army of None: Autonomous Weapons and the Future of War*, W.W. Norton & Co. at 255-257 (2018).

⁸ For a properly designed and tested autonomous system which correctly carries out the commander's intent, the commander is clearly accountable for the actions of that system. It is incumbent on states to properly design, test, and use such systems and also put in place rigorous procedures ensuring that any weapon use complies with IHL, including by ensuring individual accountability.

⁹ The Commission believes DoD's existing formulation of "appropriate human judgment," discussed in the following Judgment, captures that necessary variation and ensures that any decision to employ lethal force begins with and is under the control of human judgment, and that a human ultimately will remain accountable for any decision to employ force.

¹⁰ Press Release, U.S. Department of Defense, *DoD Adopts Ethical Principles for Artificial Intelligence* (Feb. 24, 2020), <https://www.defense.gov/Newsroom/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>.

¹¹ DoDD 5000.01 requires any weapon fielded by DoD to undergo a legal review to ensure compliance with the Law of Armed Conflict (LOAC), adhering to the requirements set out in Article 36 of the Protocol Additional to the Geneva Conventions of 12 August 1949. DoDD 3000.09 and the DoD AI Ethics Principles build on top of this baseline. See *Department of Defense Directive 5000.01: The Defense Acquisition System*, U.S. Department of Defense at 9 (Sept. 9, 2020), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/500001p.pdf?ver=2020-09-09-160307-310>; *Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I)*, 8 June 1977, International Committee of the Red Cross (last accessed Jan. 5, 2021), <https://ihl-databases.icrc.org/applic/ihl/ihl.nsf/WebART/470-750045>.

¹² *Department of Defense Directive No. 2311.01: DoD Law of War Program*, U.S. Department of Defense at 11 (July 2, 2020), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/231101p.pdf?ver=2020-07-02-143157-007>.

¹³ *A Guide to the Legal Review of New Weapons, Means and Methods of Warfare: Measures to Implement Article 36 of Additional Protocol I of 1977*, International Committee of the Red Cross at 5, n. 8 (Jan. 2006), https://www.icrc.org/en/doc/assets/files/other/icrc_002_0902.pdf.

¹⁴ *Department of Defense Directive 3000.09: Autonomy in Weapon systems*, U.S. Department of Defense at 2 (Nov. 21, 2012, incorp. change 1 May 8, 2017), <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>. The weapons-review processes established in DoDD 3000.09 are designed specifically to ensure that any U.S. autonomous weapon system complies with IHL principles such as discrimination and proportionality while also maintaining appropriate levels of human judgment and ensuring accountability.

¹⁵ *Department of Defense Instruction 5025.01: DoD Issuances Program* at 22 (Aug. 1, 2016, incorp. change 3 May 22, 2019), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/502501p.pdf?ver=2020-05-20-081854-657>.

¹⁶ See the Appendix of this report containing the abridged version of NSCAI's Key Considerations for Responsible Development & Fielding of AI. For additional details on the Commission's recommendation for future R&D needed to advance capabilities for Testing, Evaluation, Verification, and Validation of AI systems, see the section on "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁷ The DoD Law of War manual serves as a detailed resource for all DoD personnel responsible for implementing the law of war and executing military operations. See *Department of Defense Law of War Manual*, U.S. Department of Defense (Dec. 2016), <https://dod.defense.gov/Portals/1/Documents/pubs/DoD%20Law%20of%20War%20Manual%20-%20June%202015%20Updated%20Dec%202016.pdf?ver=2016-12-13-172036-190>.

¹⁸ Press Release, U.S. Department of Defense, *DoD Adopts Ethical Principles for Artificial Intelligence* (Feb. 24, 2020), <https://www.defense.gov/Newsroom/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/>.

¹⁹ David Axe, *Don't Panic, But Russia Is Training its Robot Tanks to Understand Human Speech*, *Forbes* (June 30, 2020), <https://www.forbes.com/sites/davidaxe/2020/06/30/dont-panic-but-russia-is-training-its-robot-tanks-to-understand-human-speech/?sh=7373377914f2>.

²⁰ Patrick Tucker, *SecDef: China Is Exporting Killer Robots to the Mideast*, *Defense One* (Nov. 5, 2019), <https://www.defenseone.com/technology/2019/11/secdef-china-exporting-killer-robots-mideast/161100/>.

²¹ The United States has expressed similar concerns with respect to treaties banning cluster munitions and nuclear weapons. See *Q&A: Convention on Cluster Munitions*, HRW (Nov. 6, 2010), <https://www.hrw.org/news/2010/11/06/qa-convention-cluster-munitions#>; Heather Williams, *What the Nuclear Ban Treaty Means for America's Allies*, *War on the Rocks* (Nov. 5, 2020), <https://warontherocks.com/2020/11/what-the-nuclear-ban-treaty-means-for-americas-allies/>. As of March 2021, no ally with which the United States has a mutual defense agreement has expressed support for a treaty banning LAWS.

²² The Commission recognizes that AI should assist in some aspects of the nuclear command and control apparatus, such as early warning, early launch detection, and multi-sensor fusion to validate single sensor detections and potentially eliminate false detections.

²³ *Nuclear Matters Handbook 2020*, Office of the Deputy Assistant Secretary of Defense for Nuclear Matters at 18 (2020), <https://fas.org/man/eprint/nmh2020.pdf>.

²⁴ There could be other reasons countries may delegate nuclear weapons launch authority to autonomous systems, particularly if leadership trusts machines to execute launch orders more than humans. A political agreement is unlikely to be able to address these concerns, although offering it would highlight how other nations are engaging in irresponsible and dangerous behavior.

²⁵ Michael Peck, *Russia's 'Dead Hand' Nuclear Doomsday Weapon is Back*, *The National Interest* (Dec. 12, 2018), <https://nationalinterest.org/blog/buzz/russias-dead-hand-nuclear-doomsday-weapon-back-38492>.

²⁶ Press Release, U.S. Department of State, *Deputy Secretary Sullivan's Participation in Strategic Security Dialogue with Russian Deputy Foreign Minister Sergey Ryabkov* (July 17, 2019), <https://2017-2021.state.gov/deputy-secretary-sullivans-participation-in-strategic-security-dialogue-with-russian-deputy-foreign-minister-sergey-ryabkov/index.html>; Press Release, U.S. Department of State, *The United States and Russia Hold Space Security Exchange* (July 28, 2020), <https://2017-2021.state.gov/the-united-states-and-russia-hold-space-security-exchange/index.html>.

²⁷ Over the last year, Chinese experts have participated actively in several Track II dialogues with U.S. experts on the safety of military AI systems, potentially signaling a desire for formal government-to-government communication on these issues.

Chapter 4 - Endnotes

²⁸ See Michael C. Horowitz and Paul Scharre, *AI and International Stability: Risks and Confidence-Building Measures*, Center for a New American Security (Jan. 2021), <https://s3.us-east-1.amazonaws.com/files.cnas.org/documents/AI-and-International-Stability-Risks-and-Confidence-Building-Measures.pdf?mtime=20210112103229&focal=none>.

²⁹ *Final Report of the 2019 Session of the Group of Governmental Experts on Emerging Technologies in the Area of Lethal Autonomous Weapon systems, Group of Governmental Experts of the High Contracting Parties to the Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons Which May Be Deemed to Be Excessively Injurious or to Have Indiscriminate Effect*, CCW/MSP/2019/CRP.2/Rev.1, (Nov. 13-15, 2019), <https://undocs.org/CCW/MSP/2019/9>.

³⁰ See the Appendix of this report containing the abridged version of NSCAI's Key Considerations for Responsible Development & Fielding of AI. For additional details on the Commission's recommendation for future action on International collaboration and cooperation, see the section on "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

³¹ See Chapter 14 of this report for additional information on the difficulty of using export controls to prevent the transfer of AI algorithms.

³² Along these lines, the United States shared the technology for Permissive Action Links (PALs), which prevent the unauthorized arming of a nuclear weapon, with the Soviet Union in the 1970s. It is not clear if there is an equivalent technology to PALs for AI, one which would reduce the risk of unauthorized or accidental escalation by an AI system without simultaneously significantly increasing the military performance of that system. If equivalent technologies are developed, cooperation would have to be considered on a case-by-case basis.

Chapter 5: AI and the Future of National Intelligence

2025: AI-Enabled Intelligence and Predictive Analysis



**Empowering
Science and
Technology
Leadership**



**Innovative
Approaches
to Human and
Machine Teaming**



**Capitalizing on AI
Analysis of Open-
Source Information**



**Prioritizing
the Collection
of Scientific
and Technical
Intelligence**



**Building the
IC Information
Technology
Environment**

Intelligence will benefit from rapid adoption of artificial intelligence (AI)-enabled technologies more than any other national security mission. As every possible platform—both machine and human—contributes to the global information grid, and as the number of sensors grows exponentially, the volume, velocity, and variety of data threaten to overwhelm intelligence analysis. Ascertaining the veracity and value of information will be harder. Analysts will be challenged to provide the context crucial for turning information into actionable intelligence.

AI will help intelligence professionals find needles in haystacks, connect the dots, and disrupt dangerous plots by discerning trends and discovering previously hidden or masked indications and warnings. AI-enabled capabilities will improve every stage of the intelligence cycle from tasking through collection, processing, exploitation, analysis, and dissemination. AI algorithms can sift through vast amounts of data to find patterns, detect threats, identify correlations, and make predictions. AI tools can make satellite imagery, communications signals, economic indicators, social media data, and other large sources of information more intelligible. AI can find correlations between open-source data and other sources of intelligence, and help the Intelligence Community (IC) be more precise, efficient, and effective in its targeting and collections activities. The constellation of current and emerging AI technologies applicable to intelligence missions includes computer vision for imagery analysis, biometric technologies (such as face, voice, and gait recognition), natural language processing, and algorithmic search and query functions for large databases, among others. Most important, AI enables data fusion from dissimilar data streams to create a composite picture.¹

In military scenarios—against technologically advanced adversaries, rogue states, or terrorist organizations—AI-enabled intelligence, surveillance, and reconnaissance platforms and AI-enabled indication and warning (I&W) systems will be critical for the kind of advanced warfighting capabilities discussed in Chapter 3 of this report. Through automation, AI-enabled systems will optimize tasking and collection for platforms, sensors, and assets in near-real time in response to dynamic intelligence requirements or changes in the environment. At the tactical edge, “smart” sensors will be capable of pre-processing raw intelligence and prioritizing the data to transmit and store, which will be especially helpful in degraded or low-bandwidth environments. Once collected, intelligent processing systems can triage the information, identify trends and patterns, summarize key implications, and prepare the highest-priority information for human review (or flag items of particular interest, based on analyst-defined conditions). This includes advanced I&W systems that will enable warfighters to anticipate and understand emerging threats earlier, allowing them to proactively shape the environment, as well as systems close to the tactical edge identifying adversarial denial and deception efforts. When paired with human judgment, these capabilities will enhance all-domain awareness, lead to tighter and more informed decision cycles, offer recommendations for different courses of action, and allow rapid counter-actions to adversary actions.

The need to adapt is made urgent by the quickening diffusion of these new technologies. Once exquisite IC capabilities are now in wide use around the world.² Our adversaries' ability to quickly adopt AI tools means that the IC may be more vulnerable to deception, information operations, sources and methods exposure, cyber operations, and counterintelligence activities. The IC has been an early mover within the government in establishing some of the underlying infrastructure to enable the adoption of AI, such as contracting an IC-wide commercial cloud service in 2013.³ In addition, the IC's 2019 Augmenting Intelligence using Machines (AIM) initiative provided direction and a framework for broader adoption, and some intelligence agencies have made great strides in AI adoption, putting them ahead of others in government. Still, critical barriers in authorities, policies, budgets, data sharing, and technical standards keep the IC from fully realizing its potential, and none of these recommendations will be effective without substantial reforms of the security clearance process.


An Ambitious Agenda: AI-Ready by 2025.

To build on the progress that individual agencies have made, the IC should set the ambitious goal of adopting and integrating AI-enabled capabilities across every possible aspect of the intelligence enterprise as part of a larger vision for the future of intelligence.

An AI-Ready IC by 2025:

Intelligence professionals enabled with baseline digital literacy and access to the digital infrastructure and software required for ubiquitous AI integration in each stage of the intelligence cycle.

Starting immediately, the IC should prioritize automating each stage of the intelligence cycle to the greatest extent possible and processing all available data and information through AI-enabled analytic systems before human analyst review. Products should also be disseminated at machine speed—which means they must be in machine-readable formats—and systems across the IC must be able to ingest and use them without manual intervention. Optimizing AI-enabled systems in this way will require an entirely different approach to the creation and review of finished intelligence products. The IC should require that all intelligence products include both a human-readable version and, just as important, an automated machine-readable version that can be ingested into other analytic systems throughout the IC. All future intelligence systems should be optimized for AI-oriented data collection and processing.

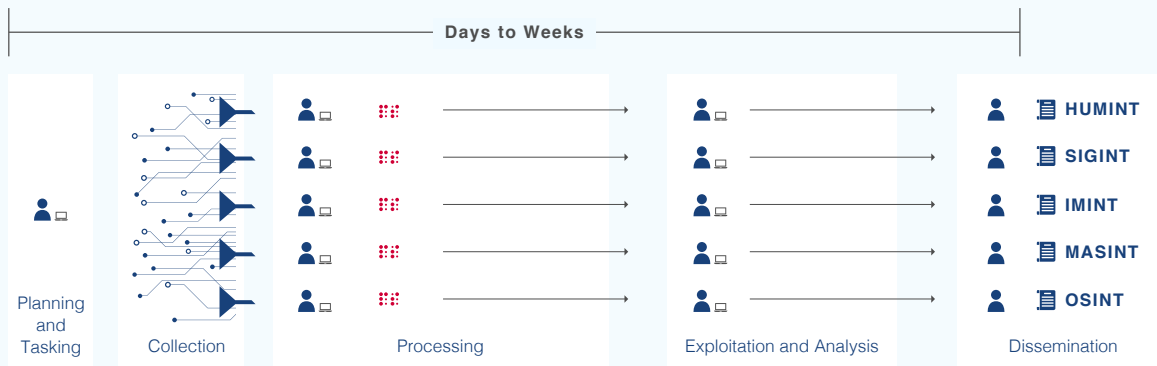


“The IC should require that all intelligence products include both a human-readable version and, as importantly, an automated machine-readable version that can be ingested into other analytic systems throughout the IC.”


Once the IC has automated its processes within individual intelligence disciplines, it should fuse those individual processes into a continuous pipeline of all-source intelligence analysis processed through a federated architecture of continually learning analytic engines. This transformational change could lead to insights arising from human-machine teaming that are beyond the current limits of unaided human cognition. Such a system would bring greater clarity to ongoing developments and also enable more accurate and reliable predictive analysis of emerging threats. As analysts gain more trust in AI-enabled systems, the ratio of human- to machine-led analysis will tip more heavily toward machines.

Current

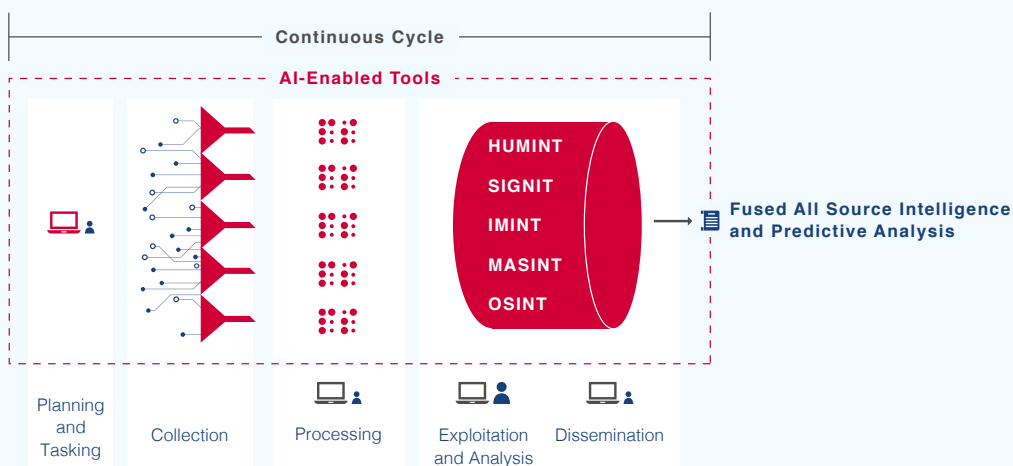
AI-Enabled National Intelligence.



Optimized: AI-Enabled Automation within Current Intelligence Disciplines



Transformed: AI-Enabled All Source Intelligence and Predictive Analysis




Preparing for an AI-ready 2025 demands the following actions:

Recommendation

Empower the IC's science and technology leadership. The Director of National Intelligence (DNI) should designate the Director of Science and Technology (S&T) within the Office of the Director of National Intelligence (ODNI) as the IC's Chief Technology Officer (CTO) and task and empower this position to drive the IC's adoption of AI-enabled applications to solve operational intelligence requirements. To do so, the IC CTO should oversee the AIM strategy, establish and enforce common technical standards and policies necessary to rapidly and responsibly scale AI-enabled applications across the IC, and lead acquisition reform to ensure that the IC can rapidly procure and field systems to its intelligence professionals. The IC CTO should be granted additional authorities for establishing policies on and supervising IC research and engineering, technology development, technology transition, appropriate prototyping activities, experimentation, and developmental testing activities.

Recommendation

Change risk management practices to accelerate new technology adoption. The IC needs to balance the technical risks involved in bringing new technologies online and quickly updating them with the substantial operational risks that result from not keeping pace, similar to DoD. Regular software upgrades should be automated to the extent possible. To share software tools more easily among agencies, reciprocal accreditation of information technology systems should be the standard.⁴



“The IC needs to balance the technical risks involved in bringing new technologies on line and quickly updating them with the substantial operational risks that result from not keeping pace ...”

To coordinate these changes, the ODNI should establish a Senior Risk Management Council focused on technology modernization.⁵ Its task should be to weigh the risks of adopting new technologies with the opportunity costs of not doing so. Its goal should be to ensure that analysts have access to the tools they need to do their jobs.

The IC will need support from the intelligence committees in Congress—for example, in the flexible use of funds within a more agile software development framework. To support the argument for greater flexibility, the IC should develop data-driven ways of communicating operational gains, as well as credible assessments of the risk of inaction.

Improve coordination and interoperability between the IC and DoD. The IC must aggressively pursue automated interoperability with the DoD for intelligence operations conducted at machine speeds.⁶ To do this, security managers and network administrators must build greater confidence in fast and secure data exchanges. ODNI, the Under Secretary of Defense for Intelligence and Security, and the Joint Artificial Intelligence Center (JAIC) should coordinate more on intelligence-related AI projects to minimize duplication of effort while maximizing common approaches to AI capability development, testing and evaluation, deployment, international engagement, and policies and authorities. They should work together to create interoperable and sharable resources and tools—such as those envisioned in the AI R&D ecosystem described in Chapter 2 of this report—and should establish a culture of sharing all AI-enabled capabilities whenever feasible.⁷

Recommendation

*Capitalize on AI-enabled analysis of open-source and publicly available information.*⁸ The IC should develop a coordinated and federated approach to applying AI-enabled applications to open-source intelligence (OSINT) and should strive to integrate open-source analysis into existing intelligence processes wherever possible in every intelligence domain.⁹

Recommendation

Prioritize and accelerate collection of scientific and technical intelligence to better understand adversary capabilities and intentions. Such collection requires the IC to significantly increase the technical sophistication, capabilities, and capacity of its analytic workforce. That must involve aggressive efforts to train, recruit, and retain analysts who have the requisite skills. These analysts must guide collection requirements and provide timely, accurate assessments. To better coordinate intelligence on these topics, including collecting on scientific and technical cooperation among our competitors, the DNI should appoint an Emerging Technology Collection Executive within the National Intelligence Council.¹⁰

Recommendation

To recruit more S&T experts into the IC, aggressively pursue security clearance reform for clearances at the Top Secret level and above, and enforce security clearance reciprocity among members of the IC. ODNI should develop and implement an AI-enabled data and science-based approach to security-clearance adjudication that significantly shortens investigation timelines.¹¹

Recommendation

Recommendation

Advance and continue to develop a purpose-built IC Information Technology Environment that can fuse intelligence from different domains and sources. An AI-enabled technical architecture of this kind could help autonomously integrate intelligence across stove-piped intelligence domains, which currently often require manual intervention to share raw data or finished analysis.¹² Doing so would help the IC blend insights from different streams of information to create a composite picture. For example, signals intelligence often depends upon human intelligence or geospatial intelligence. Likewise, the value of human intelligence can almost always be enhanced by layering signals intelligence or open-source information on top of it.

Recommendation

Embrace fused, predictive analysis as the new standard. Successfully fusing all-source/all-domain intelligence will enable accurate predictive analysis in a way that is not currently possible. The government's response to the COVID-19 virus has offered glimpses into the potential for fused data sets to inform such analysis. For example, U.S. Northern Command (working with the JAIC and the National Guard Bureau) built predictive models from dozens of different data sets that helped to identify COVID-19 hotspots and reconcile demands for critical supplies.¹³

Recommendation

Develop innovative human-centric approaches to human-machine teaming. The kind of data fusion envisioned here through autonomous machine-to-machine integration will require new concepts for human-machine teaming that optimize the strengths of each.¹⁴ The IC will need new approaches that amplify and extend human cognition to effectively handle the scale and complexity of the information generated by all-source intelligence analytic engines. When developing these systems, the IC must understand and make deliberate decisions on when and under what conditions the human or machine should act alone and under what conditions human-machine teaming is desirable.

“The kind of data fusion envisioned here through autonomous machine-to-machine integration will require new concepts for human-machine teaming that optimize the strengths of each.”

Chapter 5 - Endnotes

¹ For additional information on AI-enabled use cases throughout the intelligence cycle, see the discussion on “Applications” in *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence Through Innovation*, CSIS Technology and Intelligence Task Force at 8-22 (Jan. 13, 2021), https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/210113_Intelligence_Edge.pdf.

² *AIM Initiative: A Strategy for Augmenting Intelligence Using Machines*, Office of the Director of National Intelligence (2019), <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf> (foreword by the Honorable Sue Gordon, Principal Deputy Director of National Intelligence).

³ Frank Konkel, *The Details about the CIA’s Deal with Amazon*, *The Atlantic* (July 14, 2014), <https://www.theatlantic.com/technology/archive/2014/07/the-details-about-the-cias-deal-with-amazon/374632/>.

⁴ In adopting new software systems, the IC follows a risk-management framework developed by the National Institute of Standards and Technology (NIST). While it is a useful framework overall, it can also create delays or prevent the IC from keeping up with cutting-edge AI tools that are commercially available. For more information, see *FISMA Implementation Project*, NIST (Dec. 3, 2020), <https://csrc.nist.gov/projects/risk-management/rmf-overview>.

⁵ The Senior Risk Management Council would help the IC implement guidance from the proposed Tri-Chair Committee on Emerging Technology and function similarly to the role this commission recommended for the Under Secretary of Defense for Research and Engineering as a co-chair on the Joint Requirements Oversight Council in DoD.

⁶ For more information, see Kent Linnebur, et al., *Intelligence After Next: The Future of the IC Workplace*, MITRE Center for Technology and National Security (Nov. 1, 2020), <https://www.mitre.org/sites/default/files/publications/pr-20-1891-intelligence-after-next-the-future-of-the-ic-workplace.pdf>.

⁷ These efforts should leverage the JAIC’s Joint Common Foundation (JCF).

⁸ Pub. L. 116-260, The Consolidated Appropriations Act (2021), Division W, Section 326 (“Open source intelligence strategies and plans for the intelligence community”), Section 623 (“Independent study on open-source intelligence”), and Section 624 (“Survey on Open Source Enterprise”) provide a starting point for the IC to reimagine the role of open-source intelligence.

⁹ It is important to note that open-source intelligence (OSINT) is not limited to traditional media sources (newspapers, radio broadcasts, etc.) and social media. OSINT also includes publicly available information such as public government data sources (official reports, budget documents, hearing testimonies, etc.), professional and academic publications, commercial data sources (industry reports, financial statements, commercial imagery, etc.), and more.

¹⁰ For additional information, see the discussion on “Elevating Technical Intelligence” in *Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence Through Innovation*, CSIS Technology and Intelligence Task Force at 12 (Jan. 13, 2021), https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/210113_Intelligence_Edge.pdf.

¹¹ For more information on the need for an academic and scientific review of behavioral approaches to security clearance adjudication, see David Luckey, et al., *Assessing Continuous Evaluation Approaches for Insider Threats: How Can the Security Posture of the U.S. Departments and Agencies Be Improved?*, RAND Corporation at 28-34 (2019), https://www.rand.org/pubs/research_reports/RR2684.html.

Chapter 5 - Endnotes

¹² The technical aspects of such an environment are covered in more detail in Chapter 2 of this report.

¹³ Air Force General Terrence J. O'Shaughnessy, Commander, U.S. Northern Command & Army Lieutenant General Laura J. Richardson, Commander, U.S. Army North, *Transcript: US NORTHCOM and ARNORTH Commanders Discuss Ongoing COVID-19 Efforts*, U.S. Department of Defense (April 21, 2020), <https://www.defense.gov/Newsroom/Transcripts/Transcript/Article/2160070/us-northcom-and-arnorth-commanders-discuss-ongoing-covid-19-efforts/>.

¹⁴ See Kenneth M. Ford, et al., *Cognitive Orthoses: Toward Human-Centered AI*, *AI Magazine* at 7 (Winter 2015), <https://doi.org/10.1609/aimag.v36i4.2629>; John Laird, et al., *Future Directions in Human Machine Teaming Workshop*, U.S. Department of Defense (July 16-17, 2019), <https://basicresearch.defense.gov/Portals/61/Future%20Directions%20in%20Human%20Machine%20Teaming%20Workshop%20report%20%20%28for%20public%20release%29.pdf>; Gagan Bansal, et al., *Is the Most Accurate AI the Best Teammate? Optimizing AI for Teamwork*, *AAAI 2021* (Feb. 2021), <https://www.microsoft.com/en-us/research/publication/is-the-most-accurate-ai-the-best-teammate-optimizing-ai-for-teamwork/>.

Chapter 6: Technical Talent in Government

Improve Technical Talent in Government



Organize



Recruit



Build




Employ

The artificial intelligence (AI) competition will not be won by the side with the best technology. It will be won by the side with the best, most diverse and tech-savvy talent. The Department of Defense (DoD) and the Intelligence Community (IC) both face an alarming talent deficit. This problem is the greatest impediment to the U.S. being AI-ready by 2025. National security agencies need more digital experts now or they will remain unprepared to buy, build, and use AI and its associated technologies. Digital expertise is the most important requirement for government modernization, but few parts of government have adequately invested in building a digital workforce.¹

“DoD and the IC both face an alarming talent deficit.”

To expand its digital and AI workforce, the government needs to:



To expand its digital and AI workforce, the government needs to:

- **Organize** technologists within government through a talent management system designed to house highly skilled specialists;
- **Recruit** people who already have the skills the government needs, such as industry experts, academics, and recent college graduates;
- **Build** its own workforce by training and educating current and future government employees; and
- **Employ** its digital workforce more effectively to ensure digital talent can perform meaningful work once they are in government.




“Digital expertise is the most important requirement for government modernization ...”

The Current Model.

Government organizations responsible for creating AI solutions are struggling to build their digital workforce. Real obstacles impede recruiting and retaining AI practitioners and broader digital talent. The government does not compete with private-sector salaries and suffers from a cumbersome hiring process, and all reforms are hindered by a slow security clearance process.

We should not accept an undesirable status quo as the inevitable future. The government can compete with the private sector for talent. The government may not match private-sector salaries, but it does offer the opportunity to tackle national security challenges and to make a substantial contribution to society. The biggest obstacle hindering the recruitment of digital talent is not compensation. It is the perception, and too often the reality, that it is difficult for digital talent in government to perform meaningful work, with modern computing tools, at the forefront of a rapidly changing field.²



“We should not accept an undesirable status quo as the inevitable future. The government can compete with the private sector for talent.”

The Commission is not persuaded by the argument that the government should focus on project management and data collection and management, and outsource all development. We have heard this argument from leaders who do not believe it is feasible for the government to hire or train its own AI experts. Interestingly, we have not heard this argument from industry.



“Government strategies that do not develop a government technical workforce are short-sighted.”

Government strategies that do not develop a government technical workforce are short-sighted. Government agencies that rely solely on contractors for digital expertise will become incapable of understanding the underlying technology well enough to make successful acquisition decisions independent of contractors.³ This situation creates national security risks. While contractors should continue to play a critical role, they are incentivized, and in some sense required, to fulfill the terms of their contract, not to pursue overall system improvements or to disagree with poorly thought-out requirements or ineffective strategies. As a result, agencies that rely on contractors force their digital experts to have a secondary voice in key decisions, even those related to their field of expertise. The government will always have contractors. But the government can and should grow its own digital workforce.

Organize.

How a digital workforce is organized is as important as the workforce’s level of expertise. To generate and manage a proficient digital workforce at the scale required by the national security enterprise, the government needs to establish a talent management framework tailored to the task.

Recommendation


Departments and select agencies should create Digital Corps. We propose that departments and select agencies should create Digital Corps that would recruit, train, and educate personnel; place people in and remove them from digital workforce billets; manage digital careers; and set standards for digital workforce qualifications. Departments and select agencies would create billets for members of these Digital Corps and provide guidance to members about the work they perform for the agencies.

The Digital Corps model is inspired by the Army’s Medical Corps, which organizes experts with specialized healthcare skills that do not fit into the Army’s traditional talent management framework. Like the Medical Corps, agency-specific Digital Corps should have specialized personnel policies, guidelines for promotion, training resources, and certifications to demonstrate proficiency in new digital areas.

Recruit.

To fill these Digital Corps and to improve its broader digital workforce, the government needs to improve recruiting and the hiring process, accelerate security clearances, use temporary hiring vehicles such as the Intergovernmental Personnel Act, and build mechanisms for part-time civilian service.⁴ Many AI and other digital practitioners are interested in working with the government as either full-time employees or part-time employees. Of those desiring full-time employment, some seek an entire career as a government civilian or in the military. Others are less willing to make long-term commitments and instead desire to become temporary, full-time employees, fellows, talent exchange participants, or military reservists. A third group is willing to work with or for the government part-time, but they are unwilling to become full-time civilian employees and have no desire to serve as part of the military.

Gaps in the Recruitment Model.



Recommendation

Establish a civilian National Reserve Digital Corps. The government should tap into the pool of technologists willing to contribute part of their time to public service by creating a mechanism to hire them. While part-time employees are not a substitute for full-time employees, they can help improve AI education, perform data triage and acquisition, help guide projects and frame digital solutions, build bridges between the public and private sectors, and take on other important tasks.


To eliminate this recruitment gap, the government should establish a civilian National Reserve Digital Corps (NRDC) modeled after the military reserve's commitments and incentive structure. Members of the NRDC would become civilian special government employees in one of the agency Digital Corps and work at least 38 days each year as advisors, instructors, or developers across the government.


Recommendation

Streamline the hiring process and expand digital talent pipelines. The government hiring system's problems are well known: It moves too slowly, struggles to attract experts in a competitive market, and makes it difficult for experts who are young or do not have a degree to be hired, especially at a pay grade matching their level of expertise. These challenges are not caused by a lack of hiring authorities or an inherently slow hiring process. The Commission has been unable to identify a gap in hiring authorities for the digital workforce.

To clear this recruiting bottleneck, the government needs to expand science, technology, engineering, and mathematics (STEM) and AI talent pipelines from universities to government service, streamline the hiring process, and create agency- and military service-specific digital talent recruiting offices either for Digital Corps or agencies. The recruiting offices would monitor their corps, agency, or service's need for specific types of digital talent and be empowered to recruit technologists virtually, by attending conferences and career fairs, recruiting on college campuses, hosting prize competitions, and offering scholarships, recruiting bonuses, and referral bonuses.

National Reserve Digital Corps.






The Commission has been unable to identify a gap in hiring authorities for the digital workforce.

Standing Digital Corps will oversee government-wide progress and make recommendations to expand and improve digital talent hiring and pipelines. They should also be able to experiment with new authorities.

Build.

The government will not be able to recruit its way out of its technology workforce deficit. AI and digital talent are simply too scarce in the United States. In 2020, there were more than 430,000 open computer science jobs in the United States, while only 71,000 new computer scientists graduate from American universities each year.⁵ The government should also make a new commitment to building its workforce from the ground up with a major initiative.



“The United States needs to establish a new service academy to train future civil servants in the digital skills that are needed to modernize the government.”

Recommendation

Establish a United States Digital Service Academy. The United States needs to establish a new service academy to train future civil servants in the digital skills that are needed to modernize the government. The United States Digital Service Academy (USDSA) would be an accredited, degree-granting university that receives both government and private funding, is managed by a purpose-built independent agency within the federal government, and meets the government's needs for digital expertise—as determined by an interagency board, assisted by a Federal Advisory Committee composed of private-sector and academic technology leaders. The USDSA should be modeled off of the five U.S. military service academies but produce trained and educated government civilians for all federal government departments and agencies.

Proposed
Implementation
Plan for USDSA.

Phase One

(Years 1–2)

- Identify and secure an appropriate site for initial USDSA build-out with room for future expansion.
- Identify gaps in the government's current and envisioned digital workforce by an interagency task force under Office of Personnel Management leadership.
- Establish the USDSA administration as a new Executive Branch agency with an individual appropriation that will be responsible for the phased implementation plan and the management of the institution.
- Recruit tenure-track faculty.
- Recruit adjunct faculty, primarily from private-sector technology companies.
- Grant the USDSA the authority to accept outside funds and gifts from individuals and corporations for startup, maintenance, and infrastructure costs.
- Appropriate \$40 million to pay for administrative costs.
- Satisfy the necessary requirements set by the Department of Education, as well as the state the USDSA is in, for degree-granting approval.
- Apply for degree program-specific accreditation through the Computing Accreditation Commission on Colleges of Accreditation Board for Engineering and Technology.
- Apply for accreditation with a regional accrediting organization approved by the Department of Education and Council for Higher Education Accreditation in order to be granted "Candidate" status.
- Construct the initial physical infrastructure.
- Appropriate additional costs for the selection and purchase of the physical location and construction of the infrastructure.

Phase Two

(Years 3–5)

- Begin classes with an initial class of 500 students at the beginning of year three.
- Demonstrate compliance with all requirements and standards of the regional accrediting organization in order to be granted Membership status.

Phase Three

(Years 6–7)

- Graduate the first class.
- Ongoing improvement through accreditation assessments.
- Assess, and as appropriate, expand class sizes.



“Digitally talented people should be able to reasonably expect to spend a career performing meaningful work focused on their field of expertise in government.”

Employ.

Digitally talented people should be able to reasonably expect to spend a career performing meaningful work focused on their field of expertise in government. Without such an expectation, they are unlikely to join the government workforce, and without their experience matching expectations, they are unlikely to stay for long. Aligning expectations and experience for the digital workforce requires three changes:

- Opportunity for technologists to spend an entire career focused on the field they are passionate about;
- Well-informed leaders, some of whom are digitally proficient themselves; and
- Access to tools, data sets, and infrastructure.

These changes are more tactical than those described above, but no less impactful. Strategic initiatives succeed or fail at the tactical level, and many digital initiatives that might otherwise have strategic impact are struggling or failing tactically in part because the government does not employ its technologists effectively.

Establish new digital career fields. New career fields challenge an organization’s definition of its necessary competencies and, potentially, the nature of its identity. If the military services create career fields for software developers and data scientists, this will almost inevitably change what it means to be a soldier, sailor, airman, or marine, much as the introduction of aviation did generations ago. The government should create civilian occupational series for software development, software engineering, knowledge management, data science, and AI. The military services should create career fields in software development, data science, and AI, with both management and specialist tracks. Digital corps will need additional career fields as they develop, but these steps will establish a strong foundation.

Recommendation

Recommendation

Expand access to tools, data sets, and infrastructure. Highly skilled technologists working in government are regularly denied access to software engineering tools. The digital workforce needs access to enterprise-level software capabilities on par with those found in the private sector. Capabilities include software engineering tools, access to software libraries, vetted open-source support, curated data sets, and infrastructure for large-scale collaboration.

All career fields need improved access to the latest open-source libraries and tools.⁶ Most advanced AI and machine learning (ML) libraries need vast amounts of data available to train models on. Providing AI practitioners rich data sets across the physical and biological sciences, economics, and behavioral studies will let them focus on their areas of expertise rather than scraping obscure sources for data.

Chapter 6 - Endnotes

¹ There are pockets of excellence in several parts of the government, such as in the United States Digital Service, Kessel Run, the Army Artificial Intelligence Task Force, the USAF-MIT AI Accelerator, components of the Intelligence Community, and the national labs—but there are too few, and they have not spread widely enough across the government. Agencies' requirements for the size and type of AI workforce vary, but every agency NSCAI has engaged has expressed a need to expand its AI workforce, and the recommendations here are broadly applicable.

² NSCAI staff discussions with the Defense Innovation Board and Defense Digital Service (May 2019).

³ William A. LaPlante, *Owning the Technical Baseline*, Defense AT&L at 18-20 (July-Aug. 2015), <https://apps.dtic.mil/dtic/tr/fulltext/u2/1016084.pdf>.

⁴ For more information on the *Intergovernmental Personnel Act*, see Intergovernmental Personnel Act, OPM (last accessed Feb. 1, 2021), <https://www.opm.gov/policy-data-oversight/hiring-information/intergovernment-personnel-act/>.

⁵ Code.org (last accessed Jan. 11, 2021), <https://code.org/promote>. See also Oren Etzioni, *What Trump's Executive Order on AI Is Missing: America Needs a Special Visa Program Aimed at Attracting More AI Experts and Specialists*, Wired (Feb. 13, 2019), <https://www.wired.com/story/what-trumps-executive-order-on-ai-is-missing/>.

⁶ For the AI career field in particular, TensorFlow is one of the world's most popular libraries for training neural networks and other machine learning (ML) algorithms. PyTorch is another open-source library that aids in transforming research prototypes to production-ready machine learning models.

Chapter 7: Establishing Justified Confidence in AI Systems

Justified Confidence to Adopt and Field AI.



Leadership

Accountability and Governance



Human-AI Interaction and Teaming


Robust and Reliable AI



Testing and Evaluation, Verification and Validation (TEVV)



Artificial intelligence (AI) systems must be developed and fielded with justified confidence.¹ If AI systems do not work as designed, or are unpredictable in ways that can have significant negative consequences, then leaders will not adopt them, operators will not use them, Congress will not fund them, and the American people will not support them.



“If AI systems ... are unpredictable in ways that can have significant negative consequences, then leaders will not adopt them, operators will not use them, Congress will not fund them, and the American people will not support them.”

Achieving acceptable AI performance often is linked to the decision to accept some level of risk. No technology works perfectly under all conditions. Risk calculus changes with circumstances. The variables and considerations that inform judgments to rely on AI will vary significantly across military, intelligence, homeland security, and law enforcement missions. In a high-threat environment like combat, in some cases it may be reasonable to employ a system offering some immediate military advantage, while recognizing that it might fail; in other cases, however, a reasonable commander might want the highest assurances of AI reliability before fielding when lives are at risk.



“As departments and agencies rely more heavily on machines, a central guiding principle across national security scenarios is the continued centrality of human judgment.”

As departments and agencies rely more heavily on machines, a central guiding principle across national security scenarios is the continued centrality of human judgment. Those charged with utilizing AI need an informed understanding of risks, opportunities, and tradeoffs. They need awareness of the possibilities and limitations in a system’s expected performance. Ultimately, they need to formulate an educated answer to this question: In the given circumstance, how much confidence in the machine is enough confidence? These issues bear on the full lifecycle of an AI system—from acquisition or system development and the thresholds for justified confidence to deploy a specific AI-intensive system to the performance of the system in the field.

While there is no absolute assurance of perfection, there are policies and best practices that support making these decisions responsibly. Agencies are broadly aware of the principal challenges in employing AI systems and the necessity of incorporating best practices in the engineering and management of AI systems.

The Commission has produced a detailed framework to guide the responsible development and fielding of AI across the national security community (see the Appendix on Key Considerations for Responsible Development and Fielding of AI). It contains key considerations for policymakers and technical practitioners covering the full breadth of the AI lifecycle. The framework includes recommended practices that should be integrated and updated as the technology advances. The Commission is heartened that some departments have already taken actions to integrate recommendations from our framework.²

To assist agencies in meeting baseline criteria for Responsible AI, we highlight the main challenges and key recommendations in our framework across five issue areas.

1. Robust and Reliable AI.

Current AI systems, such as those used for perception and classification, have different kinds of failure—characterized as rates of false positives and false negatives. They are often brittle when operating at the edges of their performance competence, and it is difficult to anticipate their competence boundaries.³ They are also vulnerable to attack, and they can exhibit unwanted bias in operation. For national security missions, these can be serious problems. U.S. government agencies should:

Recommendation

Focus more federal R&D investments on advancing AI security and robustness. These investments should also advance the interpretability and explainability of AI systems, so users can better understand whether the systems are operating as intended.

Recommendation

Consult interdisciplinary groups of experts to conduct risk assessments, improve documentation practices, and build overall system architectures to limit the consequences of system failure.⁴ Such architectures should securely monitor component performance and handle errors when anomalies are detected⁵; contain AI components that are self-protecting (validating input data) and self-checking (validating data passed to the rest of the system); and include aggressive stress testing.

“The government needs AI systems that augment and complement human understanding and decision-making so that the complementary strengths of humans and AI can be leveraged as an optimal team. Achieving this remains a challenge.”

2. Human-AI Interaction and Teaming.

The government needs AI systems that augment and complement human understanding and decision-making so that the complementary strengths of humans and AI can be leveraged as an optimal team. Achieving this remains a challenge. For instance, humans are prone both to over-trusting and to under-trusting machines depending on context. Challenges also exist for measuring the performance of human-AI teams, conveying enough information while avoiding cognitive overload, enabling humans and machines to understand the circumstances in which they should pass control between each other, and maintaining appropriate human engagement to preserve situational awareness and meaningfully take action when needed. Agencies will also need to determine machine performance standards and expectations as compared with humans. The government should:

Pursue a sustained, multidisciplinary initiative through national security research labs to enhance human-AI teaming. This initiative should focus on maximizing the benefits of human-AI interaction; better measuring human performance and capabilities when working with AI systems, including testing through continuous contact and experimentation with end users; and helping AI systems better understand contextual nuances of a situation.

Recommendation


Clarify policies on human roles and functions, develop designs that optimize human-machine interaction, and provide ongoing and organization-wide AI training.

Recommendation

DoD AI Total RDT&E Investments by Research Area, FY 2015–2025

Source: Govini

- Multi-area
- Sense and Perceive
- Learn and Reason
- Hardware and Robotics
- Integrate and Assure
- Human-AI Interaction



DoD AI total RDT&E Investments.

This figure displays estimated DoD spending levels across five major research categories devised by NSCAI commissioners, indicating that investments in human-AI interaction lags behind other research categories.

Note the spending levels presented in figure represent estimates based on an analysis of DoD RDT&E budget documents for FY2021-FY2025. See *Analysis of DoD RDT&E Investments in AI*, NSCAI (on final with the Commission). Due to inherent quality issues in the source data, estimates presented contain significant, difficult to estimate margins of error.

3. Testing and Evaluation, Verification and Validation (TEVV).

Having justified confidence in AI systems requires assurances that they will perform as intended, including when interacting with humans and other systems. The TEVV of traditional legacy systems is not sufficient at providing these assurances. As a result, agencies lack common metrics to assess trustworthiness that AI systems will perform as intended. To minimize performance problems and unanticipated outcomes, an entirely new type of TEVV will be needed. This is a priority task, and a challenging one. The federal government will need to increase R&D investments to improve our understanding of how to conduct AI and software-related TEVV. Toward this end:

Recommendation

DoD should tailor and develop TEVV policies and capabilities to meet the changes needed for AI as AI-enabled systems grow in number, scope, and complexity in the Department. This should include establishing a TEVV framework and culture that integrates continuous testing; making TEVV tools and capabilities more readily available across DoD; updating or creating live, virtual, and constructive test ranges for AI-enabled systems; and restructuring the processes that underlie requirements for system design, development, and testing.⁶

Recommendation

National Institute of Standards and Technology (NIST) should provide and regularly refresh a set of standards, performance metrics, and tools for qualified confidence in AI models, data, and training environments, and predicted outcomes. NIST should lead the AI community in establishing these resources, closely engaging with experts and users from industry, academia, and government to ensure their efficacy.

4. Leadership.

Responsible development and fielding of AI requires end users and senior leaders to be aware of system capabilities and limitations so that they are not misused. It also requires subject-matter experts to support training, acquisition, risk assessment, and adoption of best practices as they evolve. Today, only the DoD has a dedicated lead for Responsible AI; employees in national security agencies taking on these roles typically do so on a voluntary, part-time basis. Without full-time dedicated staff, agencies will not succeed in fully adopting and implementing these recommended practices. The government should:

Recommendation

Appoint a full-time, senior-level Responsible AI lead in each department or agency critical to national security and each branch of the armed services. Such an official should drive Responsible AI training, provide expertise on Responsible AI policies and practices, lead interagency coordination, and shape procurement policies.

Recommendation

Create a standing body of multidisciplinary experts in the National AI Initiative Office. The standing body would provide advice to agencies as needed on responsible AI issues. The group should include people with expertise at the intersection of AI and other fields such as ethics, law, policy, economics, cognitive science, and technology, including adversarial AI techniques.

5. Accountability and Governance.

Congress and the public need to see that the government is equipped to catch and fix critical flaws in systems in time to prevent inadvertent disasters and hold humans accountable, including for misuse. Agencies need the ability to monitor AI performance as systems run (to assess if they are performing as intended) and to build systems with the necessary instrumentation to do so.⁷ Departments and agencies critical to national security and oversight entities have all expressed challenges with having visibility into their systems, while vendors are calling for clarity on instrumentation/auditability requirements. Government agencies should:

Adapt and extend existing accountability policies to cover the full lifecycle of AI systems and their components.

Recommendation

Establish policies that allow individuals to raise concerns about irresponsible AI development and institute comprehensive oversight and enforcement practices. These should include auditing and reporting requirements, a review mechanism for the most sensitive or high-risk AI systems, and appeals and grievance processes for those affected by the actions of AI systems.

Recommendation

Justified Confidence to Adopt and Field AI:

Justified confidence in AI systems will ensure AI adoption, utilization, funding, and public trust.



Leadership

Each department or agency critical to national security and each branch of the armed services should appoint a full-time, senior-level responsible AI lead. The National AI Initiative Office should create a standing body with multi-disciplinary expertise to provide advice to agencies on responsible AI issues as needed.



Accountability and Governance

Agencies should adapt and extend accountability policies to cover the full lifecycle of AI systems and their components; establish policies that allow individuals to raise concerns about irresponsible AI development; and institute comprehensive oversight and enforcement practices.



Human-AI Interaction & Teaming

National security research labs should pursue a sustained, multi-disciplinary initiative to enhance human-AI teaming. Agencies should clarify policies on human roles and functions, develop designs that optimize human-machine interaction, and provide ongoing and organization-wide AI training.



Robust and Reliable AI

The federal government should focus more federal R&D investments on advancing AI security and robustness. Agencies should consult interdisciplinary groups of experts to conduct risk assessments, improve documentation practices, and build system architectures to limit the consequences of system failure.



Testing and Evaluation, Verification and Validation (TEVV)

The federal government should increase R&D investment for T&E: The DoD should tailor and develop TEVV policies and capabilities to meet the changes needed for AI. NIST should provide and regularly refresh a set of standards, performance metrics, and tools for qualified confidence in AI models, data, and training environments, and predicted outcomes.

Chapter 7 - Endnotes

¹ The term “justified confidence,” taken from a widely used international standard, uses a specific definition of assurance as being “grounds for justified confidence.” It notes that “stakeholders need grounds for justifiable confidence prior to depending on a system” and that “the greater the degree of dependence, the greater the need for strong grounds for confidence.” *ISO/IEC/IEEE International Standard – Systems and Software Engineering – Systems and Software Assurance*, IEEE/ISO/IEC 15026-1 (2019), https://standards.ieee.org/standard/15026-1_Revision-2019.html.

² The Department of Defense’s Joint Artificial Intelligence Center (JAIC) Subcommittees on Responsible AI and Test & Evaluation have both conducted substantial mapping exercises to determine which existing practices correspond to recommendations found in the Key Considerations. Recommendations from the Key Considerations are also reinforced by inclusions in the Department of Homeland Security (DHS)’s AI Strategy. See *Department of Homeland Security Artificial Intelligence Strategy*, DHS (Dec. 2020), <https://www.dhs.gov/publication/us-department-homeland-security-artificial-intelligence-strategy?topic=intelligence-and-analysis>.

³ Like other intelligent systems, including software and humans, AI systems have competency limitations. However, we have less science to understand the performance limitations of AI systems including why, when, and how they fail.

⁴ Such interdisciplinary teams should explore the possibility of documentation/labels specifying the narrow task/mission for which a system was designed and tested. As noted in the Appendix on Key Considerations for Responsible Development & Fielding of AI, documentation of the AI lifecycle should include information about the data used to train and test a model and the methods used to test a model, both based on the context in which it will be used. It also should include requirements for re-testing, retraining, and tuning when a system is used in a different scenario or setting.

⁵ Monitoring can add a layer of robustness, but must itself also be guarded to prevent new openings for external espionage or tampering with AI systems.

⁶ Upgrades to digital infrastructure, as outlined in Chapter 2 of this report, will be required to augment physical test ranges to create digital testing environments that can leverage digital twins.

⁷ Cases in which new sensors and instrumentation are added can also introduce new vulnerabilities. It is especially important to ensure that the overall architecture of such systems is secure against external espionage and tampering.

Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security

Democratic Model of AI Governance



Invest In and Adopt AI Tools to Enhance Oversight and Auditing



Strengthen Oversight Mechanisms

Increase Public Transparency



Develop and Test Systems for Privacy and Fairness



Protect Legal Redress and Due Process



The basic purpose of the American government is to protect the security and liberty of the American people. Americans have a long tradition of debating how best to achieve these twin goals when tensions arise between them. The two decades following 9/11 saw intensive efforts to calibrate the government's powers to stop another terrorist attack with its obligations to respect individual rights and liberties. Artificial intelligence (AI) is ushering in the next era of this debate because new technologies offer government agencies more powerful ways to collect and process information, track individuals' behavior and movements, and act on the basis of computer-generated analyses.

In addition to supporting military and intelligence missions abroad, these tools are promising for national security purposes closer to home—whether to examine foreign intelligence to find signs of danger to the United States, to screen for threats at the borders, to protect against cyber attacks and information operations, or to identify domestic terrorism plots. Americans have concerns that AI applications used for these security and public safety purposes—especially those involving biometric technologies or the analysis of aggregated personal data—will invade their privacy, restrict their freedoms of speech and assembly, and reinforce bias and discrimination. At the same time, if applied effectively, AI can help improve protections for privacy and civil liberties. Machine analysis could be more precise, and AI systems potentially could enhance oversight through real-time monitoring.

For the United States, as for other democratic countries, use of AI by officials must comport with principles of limited government and individual liberty. These principles do not uphold themselves. In a democratic society, any empowerment of the state must be accompanied by wise restraints to make that power legitimate in the eyes of its citizens.

As this report argues, the promise of emerging AI technologies to enhance national security is real and significant. The ability of U.S. intelligence, homeland security, and law enforcement agencies to develop and use them for national security purposes must be preserved. To do so, however, the government must ensure that their use is effective, legitimate, and lawful. Public trust will hinge on justified assurance about compliance with privacy, civil liberties, and civil rights.

Democratic AI Governance and Novel Challenges for Privacy, Civil Liberties, and Civil Rights.

With new models of techno-authoritarian governance gaining traction abroad, the United States must continue to serve as a beacon of democratic values. The democratic model

must prove its resilience in the face of emerging technological changes that could challenge it. Fundamentally, we are confident that the American system—and the rules, norms, and institutions that uphold it—can adapt to uphold the dual imperatives of security and liberty in the AI era.

For the Intelligence Community (IC), core features of that system include laws, rules, and procedures to minimize the collection, retention, and dissemination of U.S. persons' data, as well as oversight from all three branches of government.¹ Homeland security and law enforcement agencies likewise operate within frameworks of policy, oversight, and judicial review that guide border protection and criminal investigations. Ultimately, the actions of all federal agencies are subject to the Constitution's guarantees.

Within this context, the advent of modern AI—and the novel capabilities it can bring to intelligence, homeland security, and law enforcement missions—raises a number of concerns and difficult questions and challenges with respect to the privacy, civil liberties, and civil rights of U.S. persons. For example:

- AI-powered analytics can help officials process and make sense of huge amounts of information, which can be aggregated to form a revealing “mosaic” picture of a person’s activities, whereabouts, and patterns of behavior.² Combining disparate data streams involving geolocation, web browsing, financial transactions, and other data sources creates the possibility of new insights for analysts or investigators. This could be highly useful to identify threats, but it has also raised questions about the proper scope and authorization for border or law enforcement searches.³
- Much of this personal information is held by private companies. This fact of modern digital life has raised constitutional questions about whether and when individuals should have a “reasonable expectation of privacy” in the information they provide to third parties like technology firms—and questions about the circumstances in which that information may be accessed and utilized by intelligence, homeland security, or law enforcement agencies for a legitimate national security purpose.⁴
- AI can help automate aspects of data collection and analysis. Such methods can augment the ability of analysts or investigators to sift through and triage masses of information to establish patterns or pinpoint threats. But they also raise questions about the proper roles of machine and human analysis in these processes, including for making predictive judgments. To the extent that an AI system’s functions are opaque, it may be difficult to trace and justify the computational process that led the system to make a recommendation. Determining when and how to rely on algorithms is especially pertinent to minimization and querying procedures in the IC and to building cases for law enforcement action.⁵
- AI models can evolve based on changing data and interaction with other models, leading to unexpected outcomes. As a result, AI systems require more continuous testing and evaluation than prior generations of technology.
- Unintended bias can be introduced during many stages of the machine learning (ML) process, which can lead to disparate impacts in American society, a problem that has been documented in law enforcement contexts.⁶

Tenets for Managing AI Challenges.

This Commission will not endeavor to draw all of the lines for what may be permissible or wise in particular circumstances. However, important principles to follow in different national security contexts include the following:

Foreign Intelligence Collection and Analysis: The Office of the Director of National Intelligence (ODNI) AI Ethics Guidance to the Intelligence Community is an encouraging step, as it places strong emphasis on utilizing AI for foreign intelligence missions in ways that uphold the privacy and civil liberties of Americans.⁷ As these guidelines are implemented, it will be important to pay close attention to ensuring that data minimization, retention, and querying procedures are adequate and rigorously enforced.

Border Security: AI surveillance and analysis capabilities can make the government's operations more efficient and effective at the borders and ports of entry. But to sustain public support for these uses, the Department of Homeland Security (DHS) must take care to ensure that automated screening processes lead agents only to the information they need and are authorized to access, and do not impermissibly single out individuals based on characteristics such as race or religion.

Domestic Security and Public Safety: Rapid advances in AI-enabled technologies for law enforcement purposes, including biometric surveillance techniques such as facial recognition, may be outpacing rules for their proper use. The government must exercise special caution in managing risks to bedrock constitutional principles including equal protection, due process, freedom from unreasonable searches and seizures, and freedoms of speech and assembly.⁸

In carrying out these missions, it will be important to maintain clear distinctions between appropriate authorities in these different national security contexts. It is also important to gain greater public confidence by enhancing transparency, improving the performance and reliability of AI technologies, ensuring due process, and strengthening oversight. With these tenets in mind, the government should take the following steps.

“Agencies should assess near-term opportunities and research gaps in applications of AI to address privacy and civil liberties challenges ...”

Recommendations.

Invest in and adopt AI tools to enhance oversight and auditing in support of privacy and civil liberties. Agencies should assess near-term opportunities and research gaps in applications of AI to address privacy and civil liberties challenges, such as ML techniques for classification, recommendation, anomaly detection, and other applications.⁹ Examples of advances in AI to improve auditing include tools that support financial audits and model risk management. Agencies should examine the utility of these and other current or emerging practices.¹⁰

Recommendation

Improve public transparency about how the government uses AI. There is a lack of transparency into agency policies and procedures and into the accuracy of AI systems that may impact civil liberties.¹¹ The “black box” nature of some ML systems only adds to this opacity.¹² More transparency could help to ease public concerns. Of course, in certain operational contexts, especially for intelligence and law enforcement agencies, secrecy is essential to the mission. However, existing transparency mechanisms could be utilized more effectively, and, in some cases, revised. New agency reporting requirements would also be beneficial.

Recommendation

- For AI systems that impact U.S. persons, Congress should require AI Risk Assessment Reports and AI Impact Assessments from the Intelligence Community, the Department of Homeland Security, and the Federal Bureau of Investigation. These should assess the privacy, civil liberties and civil rights implications for each qualifying AI system or significant system refresh.¹³
- DHS and the FBI should also improve practices for issuing system of records notices and privacy impact assessments to provide a more holistic view of the role of AI systems before they are fielded.

Recommendation

Develop and test systems with the goal of advancing privacy preservation and fairness. ML systems in particular require ongoing assessments of privacy and fairness assurances, including the specific definition of fairness being assumed. Although an ML system may meet requirements at a static point in time, ongoing compliance is not a given once the system is operational. This is in large part due to changing data, the introduction of unintended bias, and potential re-identification of anonymized data.¹⁴ This is a complex technical area, and continued work in the technical, legal, and policy domains is required to find greater consensus on technical approaches to preserving privacy, civil liberties, and civil rights.¹⁵ Meanwhile, agencies should take several steps:

- **Assess and mitigate risks in the design, development, and testing of AI systems.** In addition to conducting risk assessments for the privacy, civil liberties and civil rights of U.S. persons, IC elements, DHS, and the FBI should take measures to mitigate those risks, and document remaining risks that are accepted. In doing so, they should adopt practices from the Key Considerations, including using privacy protections such as robust anonymization, and when possible, privacy-preserving technology; taking steps to mitigate bias in development and testing; and assessing model performance on an ongoing basis.¹⁶
- **Identify an office, committee, or team in each agency that will conduct a pre-deployment review of AI technologies that will impact privacy, civil liberties, and civil rights.** This should include review in advance of their deployment and for compliance over the lifespan of the system. An office in each IC element, DHS, and the FBI should be equipped to assess data, model, and system documentation, and to assess the testing results of systems with respect to their intended use.
- **Establish third-party testing centers for national security-related AI systems that could impact U.S. persons.** Such independent, third-party testing could be done by a national laboratory, a University Affiliated Research Center, or a Federally Funded Research and Development Center. Such testing should be mandatory for high-stakes systems but otherwise voluntary.¹⁷ It would provide agencies with additional expertise to help overcome in-house limitations.

Recommendation

Strengthen the ability of those impacted by government actions involving AI to seek redress and have due process. AI systems will make errors.¹⁸ Agencies have to accept non-zero false positive and false negative rates in order to deploy any AI system. It is important to ensure opportunities for redress, consistent with the constitutional principle of due process—for example, when a system error leads to a benefit being denied (e.g., visa approval); restrictions on movement (e.g., being placed on a no-fly list); or an arrest.¹⁹ There are also due process concerns in cases in which AI contributes to building a case to press criminal charges.²⁰ We recommend two steps to start addressing these issues:

- **Review DHS and FBI policies and practices that may impact due process and the ability to seek redress.** DHS and the FBI should review agency policies and practices to ensure that parties aggrieved by government action involving AI technology, including through system actions or misuse, can seek redress and clearly know how to do so. This review should include whether adequate notice of AI use in decision-making is provided to impacted parties, as well as the degree to which AI systems can be audited to trace the process by which a system arrived at a recommendation, if it is contested.

- **Issue Attorney General guidance on AI and due process.** The guidance should describe how relevant agencies should safeguard the due process rights of U.S. persons when AI use may lead to a deprivation of life or liberty.

Strengthen oversight mechanisms to address current and evolving concerns. The advancement of AI requires a forward-looking approach to oversight that anticipates the continued evolution and adoption of new technologies and better positions the government to responsibly manage their employment well into the future.

Recommendation

“The advancement of AI requires a forward-looking approach to oversight that anticipates the continued evolution and adoption of new technologies, and better positions the government to responsibly manage their employment well into the future.”

The government should:

- **Establish a task force to assess the privacy and civil liberties implications of AI and emerging technologies.** The goal of the task force would be to identify gaps and make recommendations to ensure that uses of AI and associated data in U.S. government operations comport with U.S. law and values, and to study organizational reforms that would support this goal. Specifically, it should assess existing policy and legal gaps for current AI applications and emerging technologies, and make recommendations for:
 - o legislative and regulatory reforms on the development and use of AI and emerging technologies and associated data²¹; and
 - o institutional changes to ensure sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies.

- **Strengthen the ability of the Privacy and Civil Liberties Oversight Board (PCLOB) to provide meaningful oversight and advice on AI use for national security.** Since its creation in 2007, following a recommendation of the 9/11 Commission, PCLOB has had an especially important role in overseeing, and advising the government on, U.S. counterterrorism missions. In recent years, it has started turning attention to the use of new technologies in foreign intelligence collection and analysis.²² The board should be given visibility into AI systems before they are fielded, including at a more granular technical level, and should be resourced and staffed to fulfill the more technically sophisticated mission that the AI era now requires.²³
- **Empower DHS Offices of Privacy and Civil Rights and Civil Liberties (CRCL).** The Chief CRCL Officer, in coordination with the Privacy Officer, must play an integral role in the legal and approval processes for the procurement and use of AI-enabled systems, including for associated data used in DHS ML systems.
- **Require stronger coordination and alignment among federal oversight and audit organizations.** Compliance by agencies with AI documentation and testing requirements should be supported by rigorous, technically informed oversight. To achieve this and overcome current audit and oversight impediments, a standing body should align and coordinate to enhance AI oversight and audit with respect to privacy, civil liberties, and civil rights.²⁴

Democratic Model of AI Governance

AI governance poised to meet current and evolving needs for ensuring privacy, civil liberties, and civil rights in AI use for national security.



Invest in and adopt AI tools to enhance oversight and auditing.

Agencies should assess near-term opportunities and research gaps in applications of AI to address privacy and civil liberties challenges, and examine advances in AI tools to improve auditing.



Increase public transparency.

Congress should require AI Risk Assessment Reports and AI Impact Assessments from the IC, DHS, and FBI for each qualifying AI system or significant system refresh. DHS and the FBI should improve practices for issuing system-of-record notices and privacy impact assessments to provide a holistic view of the role of AI systems before they are fielded.



Develop and test systems for privacy and fairness.

Agencies should assess and mitigate risks in the design, development, and testing of AI systems; and identify an office, committee, or team in each agency to conduct a pre-deployment review of AI technologies that will impact privacy, civil liberties and civil rights. Congress should establish third-party testing centers for national security-related AI systems that could impact U.S. persons.



Strengthen oversight mechanisms.

The government should establish a task force to assess privacy and civil liberties implications of AI and emerging technologies; strengthen the Privacy and Civil Liberties Oversight Board's ability to provide oversight and advice on AI use for national security; empower DHS Offices of Privacy and Civil Rights and Civil Liberties; and require stronger coordination among federal oversight and audit organizations.



Protect legal redress and due process.

DHS and the FBI should review agency policies and practices that may impact due process and the ability to seek redress. The Attorney General should issue guidance on AI and due process.

Chapter 8 - Endnotes

¹ For a compilation of Attorney General guidelines from the IC components, see *Status of Attorney General Approved U.S. Person Procedure under E.O. 12333*, ODNI (July 14, 2016), https://www.dni.gov/files/documents/Table_of_EO12333_AG_Guidelines%20for%20PCLOB_%20Updated%20July_2016.pdf. Elements of the IC oversight system include counsels and privacy officials within intelligence agencies, the Department of Justice, independent bodies such as the Privacy and Civil Liberties Oversight Board, Federal courts including the Foreign Intelligence Surveillance Court, and the House and Senate intelligence committees.

² On the mosaic concept, see, e.g., Steven M. Bellovin, et al., *When Enough Is Enough: Location Tracking, Mosaic Theory, and Machine Learning*, NYU Journal of Law & Liberty, Vol. 8 (Sept. 3, 2013), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2320019.

³ For an informative discussion of evolving debates over Fourth Amendment regulation of government searches in the context of AI, see James E. Baker, *The Centaur's Dilemma: National Security Law for the Coming AI Revolution*, Ch. 6 (Brookings, 2020).

⁴ Congress and the Judiciary will need to assess the adequacy of current legal constraints over the federal government's obtainment and use of third-party data, including data acquired from data brokers. Either through evolving case law or legislation, agencies would benefit from clarity surrounding the Fourth Amendment's application with respect to third-party data. On third-party doctrine, see Richard M. Thompson II, *The Fourth Amendment Third-Party Doctrine*, Congressional Research Service (June 5, 2014), <https://fas.org/sgp/crs/misc/R43586.pdf>.

⁵ For a discussion and different views on the implications of human and machine analysis in the intelligence context, see Robert Litt, *The Fourth Amendment in the Information Age*, Yale Law Journal (April 27, 2016), <https://www.yalelawjournal.org/forum/fourth-amendment-information-age>; Cindy Cohn, *Protecting the Fourth Amendment in the Information Age: A Response to Robert Litt*, Yale Law Journal (July 27, 2016), <https://www.yalelawjournal.org/forum/protecting-the-fourth-amendment-in-the-information-age>.

⁶ Concerns about algorithmic error rates and disparate performance across age, skin tones, and genders are especially pronounced for facial recognition. See Patrick Grother, et al., *Face Recognition Vendor Test, Part 3: Demographic Effects*, NIST (Dec. 2019), <https://nvlpubs.nist.gov/nistpubs/ir/2019/NIST.IR.8280.pdf>. The Gender Shades Project found that various facial recognition systems were very accurate for white men, but they were significantly less accurate for women and people of color (and worst for women of color). See Gender Shades (last accessed Jan. 11, 2021), <http://gendershades.org/>.

⁷ See *Principles of Artificial Intelligence Ethics for the Intelligence Community*, ODNI (last accessed Jan. 11, 2021), <https://www.dni.gov/index.php/features/2763-principles-of-artificial-intelligence-ethics-for-the-intelligence-community>.

⁸ Some observers have found a "chilling effect" that impacts the degree to which individuals exercise freedoms of expression, association, and assembly. See, e.g., Rachel Levinson-Waldman, *Hiding in Plain Sight: A Fourth Amendment Framework for Analyzing Government Surveillance in Public*, Emory Law Journal Vol. 66 (2017), <https://scholarlycommons.law.emory.edu/elj/vol66/iss3/4/>.

⁹ Xuning (Mike) Tang & Yihua Astle, *The Impact of Deep Learning on Anomaly Detection*, Law.com (Aug. 10, 2020), <https://www.law.com/legaltechnews/2020/08/10/the-impact-of-deep-learning-on-anomaly-detection/>.

¹⁰ See, e.g., Bernhard Babel, et al., *Derisking Machine Learning and Artificial Intelligence*, McKinsey & Company (Feb. 19, 2019), <https://www.mckinsey.com/business-functions/risk/our-insights/derisking-machine-learning-and-artificial-intelligence>; Saqib Aziz & Michael Dowling, *Machine Learning and AI for Risk Management*, *Disrupting Finance* at 33-50 (Dec. 7, 2018), https://link.springer.com/chapter/10.1007/978-3-030-02330-0_3.

¹¹ For instance, disclosure by U.S. Customs and Border Protection (CBP) when using facial recognition at airports has been inconsistent, and claims exist that the FBI failed to provide information about its Next Generation Identification database and use of facial recognition as required by law. In 2020, the U.S. Government Accountability Office (GAO) found that “CPB’s privacy notices—which inform the public about its use of this technology—were not always current or available [at airports] where this technology is being used or on CBP’s website.” *Facial Recognition: CBP and TSA Are Taking Steps to Implement Programs, but CBP Should Address Privacy and System Performance Issues*, GAO (Sept. 2, 2020), <https://www.gao.gov/products/GAO-20-568>; see also *The Perpetual Line-Up: Unregulated Police Face Recognition in America*, Georgetown Law Center on Privacy & Technology (Oct. 18, 2016), <https://www.perpetuallineup.org/>.

¹² In a 2018 report, GAO has raised concerns about lack of transparency from tech companies that build algorithms and “limited testing on the systems for accuracy.” *Artificial Intelligence: Emerging Opportunities, Challenges, and Implications*, GAO (March 2018), <https://www.gao.gov/assets/700/690910.pdf>.

¹³ The Commission proposes that the task force described in this chapter, and in the accompanying Blueprint for Action, should provide binding guidance on two issues: first, when the IC, DHS, and FBI should prepare and publish AI Risk Assessment Reports and AI Impact Statements; and second, what should constitute a “qualifying AI system or significant system refresh.”

¹⁴ For example, pseudonymized data can be linked with other data to uncover a cell phone owner’s identity. See Byron Tau & Michelle Hackman, *Federal Agencies Use Cellphone Location Data for Immigration Enforcement*, Wall Street Journal (Feb. 7, 2020), <https://www.wsj.com/articles/federal-agencies-use-cellphone-location-data-for-immigration-enforcement-11581078600?mod=breakingnews>.

¹⁵ See the Appendix of this report containing the abridged version of NSCAI’s Key Considerations for Responsible Development & Fielding of AI. For further discussion of recommendations to: (1) employ technologies and operational policies that align with privacy preservation and mitigate unwanted bias and (2) to continuously monitor and evaluate AI system performance, see sections “Aligning Systems and Uses with American Values and the Rule of Law” and “System Performance” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁶ *Key Considerations for Responsible Development & Fielding of Artificial Intelligence*, NSCAI (July 2020), <https://www.nscai.gov/previous-reports/>.

¹⁷ To provide agencies guidance on when such a test mechanism should be leveraged, an organization should establish guidance on thresholds by which agencies would be required to conduct third-party testing. This should include criteria for when an AI system may pose high enough risk for privacy, civil liberties, and civil rights that it would trigger a testing requirement by a third-party auditor.

¹⁸ See, e.g., Kashmir Hill, *Wrongfully Accused by an Algorithm*, New York Times (June 24, 2020), <https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html>.

¹⁹ An individual’s right to due process, including notice, is grounded in the Constitution and the case law expounding on that right. Our recommendation’s aim is to help ensure that the government does its part to uphold these long-standing rights notwithstanding its use of AI.

Chapter 8 - Endnotes

²⁰ Due process rights require that individuals have the ability to meaningfully challenge a decision made against them. In federal criminal trials, this includes the government explaining how an unfavorable decision was reached, so it can be contested. In cases where AI-assisted or AI-enabled decisions are made, certain AI techniques will be less conducive to due process. See Danielle Keats Citron, *Technological Due Process*, Washington University Law Review (2008), https://openscholarship.wustl.edu/cgi/viewcontent.cgi?article=1166&context=law_lawreview; Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, Emory Law Journal (March 9, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3553590. Early cases in which an AI system's predictions, classifications, or recommendations have been challenged in court illustrate that defendants encounter substantial impediments in seeking to exercise their rights. See *Litigating Algorithms: Challenging Government Use of Algorithmic Decision Systems*, AI Now Institute (Sept. 2018), <https://ainowinstitute.org/litigatingalgorithms.pdf>. There are also open questions including federal rules of evidence and criminal procedure as they relate to AI. For instance, evidentiary standards for admitting AI evidence in court have yet to be developed.

²¹ Examples include baseline AI standards and policy guidance for biometric identification technologies, for government procurement of commercial AI products, and for federal data privacy standards.

²² See *Projects*, PCLOB (last accessed Jan. 9, 2021), <https://www.pcllob.gov/Projects>.

²³ PCLOB works alongside multiple oversight organizations to conduct oversight. It will also be important for PCLOB and such organizations to better align and coordinate to conduct complementary AI oversight and auditing with respect to privacy, civil liberties, and civil rights.

²⁴ For examples of impediments, see Taka Ariga & Stephen Sanford, *A is for Accountability: Oversight in the Age of Artificial Intelligence*, ECA Journal at 88-91 (Jan. 2020), https://www.eca.europa.eu/Lists/ECADocuments/JOURNAL20_01/JOURNAL20_01.pdf; see also Press Release, Office of the Inspector General of the Intelligence Community, *The Inspector General of the Intelligence Community Issues Statement on Artificial Intelligence* (May 30, 2019), <https://www.dni.gov/files/ICIG/Documents/News/ICIG%20News/2019/May%2030%20-%20AI/Press%20Release%20-%20AI.pdf>; Michael K. Atkinson, *Semiannual Report: October 2018–March 2019*, Office of the Inspector General of the Intelligence Community (2019), https://www.oversight.gov/sites/default/files/oiig-sa-reports/20190430_ICIG-SAR_Oct18-Mar19.pdf.

PART TWO



PART II: WINNING THE TECHNOLOGY COMPETITION	155
Chapter 9: A Strategy for Competition and Cooperation	157
Chapter 10: The Talent Competition	171
Chapter 11: Accelerating AI Innovation	183
Chapter 12: Intellectual Property	199
Chapter 13: Microelectronics	211
Chapter 14: Technology Protection	223
Chapter 15: A Favorable International Technology Order	241
Chapter 16: Associated Technologies	253

Chapter 9: A Strategy for Competition and Cooperation

Organizing the U.S. Government to Tackle Emerging Technology Challenges



The impact of artificial intelligence (AI) on the world will extend far beyond narrow national security applications. The development of AI constitutes a new pillar of strategic competition, and it heightens the competition in existing pillars. The nation with the most resilient and productive economic base will be best positioned to seize the mantle of world leadership. That base increasingly depends on the strength of the innovation economy, which in turn will depend on AI. AI technologies will drive waves of advancement in critical infrastructure, commerce, transportation, health, education, financial markets, food production, and environmental sustainability.

The race to research, develop, and deploy AI and associated technologies is already intensifying strategic competition. The U.S. government must embrace the AI competition and organize to win it. The American approach to innovation, which has served the country well for decades, must be recalibrated to account for the centrality of the competition involving AI and associated technologies to the emerging U.S.-China rivalry. To retain its innovation leadership and position in the world, the United States needs a stronger government-led technology strategy that integrates promotion and protection policies and links investments in AI to a larger constellation of related emerging technologies.¹




“The race to research, develop and deploy AI and associated technologies is already intensifying strategic competition. The U.S. government must embrace the AI competition and organize to win it.”

This chapter articulates the nature of the AI competition and the two prerequisites for winning it: organizing for technology competition under White House leadership and establishing the principles for continued cooperation with competitors. The following chapters (10–16) enumerate the core elements of an integrated strategy and prescribe actions to ensure the United States wins the AI competition and sets the foundation to win the broader technology competition. It is foremost an affirmative agenda for growing and recruiting talent, promoting a diverse AI innovation ecosystem, and investing in the R&D to harness AI and associated technologies to build a healthier, more prosperous, and secure society. Protection of research, intellectual property (IP), and investments will be necessary to complement the effort to invigorate AI competitiveness at home and build a coalition of like-minded partners in the world.

The United States Government must understand and define the technology competition, organize for it, and set the terms to engage with China.

Made in China 2025 and AI World Leader 2030:

China has already developed a strategy for technology leadership, picked key technology sectors, started high-tech projects within key sectors, and delegated authority across individual government agencies.



Organizing for the Competition

Understanding the Competition.

Establish a White House National Technology Competitiveness Council:

Empower a single entity in the White House to set strategic direction and oversee a coordinated approach to technology competition.

Managing the Competition

Begin a U.S.-China Comprehensive S&T Dialogue:

Establish a high-level dialogue with China to discuss challenges and manage tensions related to emerging technologies (e.g. AI, quantum, biotech).

Winning the Competition

Develop a National Technology Strategy:

The Technology Competitiveness Council should develop a national strategy to guide U.S. policy across the constellation of emerging technologies starting with AI.

The U.S.-China AI Competition Is Serious and Complex.

The leading indexes that measure progress in AI development generally place the United States ahead of China.² However, the gap is closing quickly. China stands a reasonable chance of overtaking the United States as the leading center of AI innovation in the coming decade.³ In recent years, technology firms in China have produced pathfinding advances in natural language processing,⁴ facial recognition technology,⁵ and other AI-enabled domains. China's businesses, investors, technologists, and academics are integral to global AI development. China's social media and e-commerce companies compete for users around the world. Its telecoms build global 5G infrastructure. Its venture capitalists and large technology firms invest huge sums in new startups.⁶ Its leading AI companies have research labs in the United States⁷ and elsewhere.⁸ Its researchers produce a trove of respected AI papers that advance the field.⁹ None of this would concern us from a national security perspective, except for the fact that China is led by a single-party authoritarian regime that threatens American interests.



“China stands a reasonable chance of overtaking the United States as the leading center of AI innovation in the coming decade.”

China has moved more quickly and with more determination than the United States, guided by a constellation of AI plans for government ministries, universities, and companies.¹⁰ These strategic documents reflect Beijing's view that advances in AI will fundamentally reshape military and economic competition in the coming decades.¹¹ China has backed up its strategic plans with significant state subsidies to technology firms and academic institutions that engage in cutting-edge AI research.¹² China preserves its capital by taking advantage of basic research done by the West so that it can focus more on applications. It pours significant sums of money into research and talent in relevant fields,¹³ and it promotes “national champion” companies to win markets abroad.¹⁴ Through its military-civil fusion programs, China has sought to integrate advances in AI from the commercial and academic worlds into military power.¹⁵ Using espionage, technology transfer programs, and targeted investment, Beijing seeks to acquire sensitive IP and data from the United States and other countries.¹⁶

The U.S.-China competition is complicated by the complex web of supply chains, research partnerships, and business relationships that link the world's two AI leaders. Dramatic steps to sever these ties could be costly for Americans and reverberate across the world. The relationships between American and Chinese academics, innovators, and markets are deep, often mutually beneficial, and help advance the field of AI.¹⁷ Moreover, it remains in the U.S. national interest to leverage formal diplomatic dialogue about AI and other emerging technologies and to explore areas for cooperative AI projects that will benefit humanity.



“The United States can compete against China without ending collaborative AI research and severing all technology commerce.”

The United States can compete against China without ending collaborative AI research and severing all technology commerce. Broad-based technological decoupling with China could deprive U.S. universities and companies of scarce AI and science, technology, engineering, and mathematics (STEM) talent,¹⁸ sever American companies' efficient supply chains,¹⁹ and cut off access to markets and capital for innovative firms.²⁰ Instead, the United States should conceive of targeted disentanglement as just one element of its overall approach, which, if applied judiciously to key sectors, can help build U.S. technological resilience, reduce threats from illicit technology transfer, and protect national security—critical sectors.

The Policy Challenges.

China's competitive approach should not define the U.S. approach to innovation, but it does present an alternative model of AI development, frame the stakes of competition, and expose the sheer breadth of public policy choices the U.S. government must make to preserve American advantages.

The United States will need to reexamine its immigration policies to ensure that America wins the competition for AI talent. It will need to consider AI and broader STEM education initiatives through the lens of global competitiveness. It will have to consider how to diversify the AI research agenda and expand access to the data and tools necessary to

conduct AI research in the face of costs for compute and data that are consolidating AI in fewer locations and shifting the balance from universities to the private sector. The United States will have to consider whether the long-standing approaches to IP are best suited to an era in which IP theft is pervasive and the U.S. IP regime has not yet fully accounted for AI and other emerging technologies. The United States will need to protect its leadership in the design of microelectronics, which may include encouraging the domestic reshoring of critical manufacturing on which our national security depends. And the United States will have to ensure that its tools and policies designed to prevent illicit technology transfer are postured to address the national security challenges presented by dual-use emerging technologies.

These AI-specific challenges, in turn, expose even more fundamental questions spanning the technology, economic, and national security spheres:

- How to compete with a rival without compromising U.S. values—including free-market principles, individual liberty, and limited government.
- How to ensure the proper balance between defense and economic priorities.
- How to preserve hardware advantages without suffocating the domestic designers and producers that rely on foreign competitors' markets.
- How to capitalize on and shape private-sector developments for national security ends without stifling private sector-led and free-market innovation.
- How to draw on the best global talent without enabling damaging technology and knowledge transfer to competitors.
- How to foster an open collaborative research environment while closing licit and illicit loopholes exploited by foreign competitors.
- How to sustain long-term strategies for R&D that are nevertheless responsive to rapidly shifting geopolitical and technology developments.
- How to ensure the free flow of investment/capital without allowing strategic competitors to buy strategic advantage.
- How to engage with our allies and other partners to reduce their dependence on China's digital technologies, build more resilient supply chains, and develop technology standards and norms that reflect democratic values.

The Need for a Stronger Government Role in Technology Strategy.

The Commission is not calling for a state-directed economy, a five-year plan, or China-style “military-civil fusion.” It is instead urging a government-led process to restore a more balanced equilibrium between government, industry, and academia that ensures a

diverse research environment, competitive economy, and the sustainment of a research agenda that supports the needs of the nation. The government has a long history of mobilizing industry and academia and making huge investments when the United States is challenged.²¹ Against the backdrop of a declared and committed competitor like China, and given AI's transformative potential, the United States is confronting such a moment.



“... the U.S. government champions AI leadership in speeches and memorandums, but deploys few resources relative to commercial investment and historic funding benchmarks ...”

Today, the U.S. government champions AI leadership in speeches and memorandums, but it deploys few resources relative to commercial investment and historic funding benchmarks and relies on a decentralized governance structure for achieving it.²² There is talk of a global talent competition, but in recent years the United States has tightened restrictions on visas for highly skilled workers,²³ and U.S. students at the kindergarten to 12th grade (K-12) level have lagged behind East Asian and European competitors in exams designed to measure competency in STEM fields.²⁴ Tech leaders and government officials talk about the importance of “public-private partnership,” but there is little action in either direction to deepen it in concrete ways. U.S. experts warn of the danger of AI being used for techno-authoritarian ends,²⁵ but Washington has not led any new enduring coalition to create democratic alternatives. Current policies amount to a compilation of disparate AI-related activities underway in the federal government. Nowhere can one find a strategy coupled with the organization and resources to win an AI competition and preserve the United States' AI leadership.

The government will have to orchestrate policies to promote innovation; protect industries and sectors critical to national security; recruit and train talent; incentivize domestic research, development, and production across a range of technologies deemed essential for national security and economic prosperity; and marshal coalitions of allies and

partners to support democratic norms. Some elements of a national strategy will need to be coordinated and replicated at the state level, through state-specific strategies to support AI research, commerce, and education. This will require a complex sequencing of promotion and protection actions to minimize costs and risks of punitive actions; ensure basic and applied research agendas are mutually reinforcing; coordinate approaches with international partners; and align executive priorities with legislative powers. It will require identifying technology trends and assessing the relative strengths of the United States and its competitors. It will require, above all, strong and consistent White House leadership.

The Case for White House Leadership.

The government will require a center of power that can exert gravitational pull on domestic economic, national security, and science and technology policies. We have no such organization today. Several separate Executive Office of the President (EOP) entities possess some responsibility and capacity to fulfill the basic organizational requirements: the National Security Council (NSC),²⁶ the Office of Science and Technology Policy (OSTP)²⁷ and its associated National Science and Technology Council (NSTC),²⁸ and the National Economic Council (NEC).²⁹ The Domestic Policy Council (DPC) also has critical related responsibilities and a similar mandate with leadership in the realm of immigration policy, education policy, and regulatory policy.³⁰ An additional entity—the Office of Management and Budget (OMB)—oversees related budgets and government reform efforts.

In the absence of an overarching structure, it is left to the President and Vice President to identify, adjudicate, and reconcile the positions that emerge from parallel interagency processes, while leaving endless room for gadflies to run the gaps and influence the President. The President needs a tool to help decide and drive a new technology strategy down through the necessary but not sufficient existing councils and into the rest of the government. The White House should:

Technology
Competitive
Council.



“The government will require a center of power that can exert gravitational pull on domestic economic, national security, and science and technology policies. We have no such organization today.”

Create a Technology Competitiveness Council. The United States must strengthen executive leadership in technology policy in the White House by empowering a single entity to implement a comprehensive technology strategy. The Commission proposes creating a new Technology Competitiveness Council (TCC), which would include the same amalgamation of EOP leaders and Cabinet secretaries as other White House forums for convening the interagency, and be chaired by the Vice President with a newly appointed Assistant to the President for Technology Competitiveness serving as the day-to-day leader. The TCC would ensure that the gaps between NEC, OSTP, and NSC responsibilities are filled and linked to OMB. It would not replace the NSC, NEC, or OSTP-led NSTC structures, but would provide a forum for reconciling competing security, economic, and scientific priorities and elevate technology policy and concerns from a technical to a strategic level. To coordinate the council’s work, it is necessary to create a new principal, the Assistant to the President for Technology Competitiveness, responsible for ensuring policies pertaining to emerging technologies receive sufficient Presidential-level attention.

Recommendation

Develop a National Technology Strategy. The TCC should create a National Technology Strategy, building on the elements we present here, which can guide U.S. policy across all key emerging technologies starting with AI. The goal of the National Technology Strategy should be to ensure long-term, overall U.S. leadership in technology, particularly emerging technologies critical to national security and competitiveness. The strategy should weigh the difficult tradeoffs between competing policy interests and priorities, identify critical technologies where competitors have sought to match or overtake U.S. leadership, and facilitate an integrated policy approach to emerging technologies. As a starting point, the strategy should build on the following pillars: 1) winning the AI talent competition; 2) promoting American AI innovation; 3) protecting U.S. AI advantages; and 4) leading a favorable international AI order.

Recommendation

Recommendation

Establish a high-level U.S.-China Comprehensive Science & Technology dialogue. The United States should establish a regular, high-level diplomatic dialogue with China that benefits the American people, remains faithful to our allies, and presses China to abide by international norms. The dialogue should focus on challenges presented by emerging technologies—to include AI, biotechnology, and other technologies as agreed by both sides. The dialogue should have two overarching objectives:

- Identify targeted areas of cooperation on emerging technologies to solve global challenges such as climate change and natural disaster relief; and
- Provide a forum to air a discrete set of concerns around specific uses of emerging technologies while building relationships and establishing processes between the two nations.

Chapter 9 - Endnotes

¹ While the U.S. government has released a number of documents emphasizing the importance of AI research and development—see, for example, President Trump’s executive order on AI—the U.S. lacks a comprehensive, whole-of-government plan to guide policymakers, researchers, and businesses toward a more secure U.S. future. *Artificial Intelligence for the American People*, The White House (last accessed Jan. 28, 2021), <https://trumpwhitehouse.archives.gov/ai/>.

² See, e.g., Alexandra Mousavizadeh, et al., *The Global AI Index*, Tortoise Media (Dec. 3, 2019), <https://www.tortoisemedia.com/2019/12/03/global-ai-index/>; Jean François Gagné, et al., *Global AI Talent Report 2020* (last accessed Dec. 29, 2020), <https://jfgagne.ai/global-ai-talent-report-2020/>; *The Global AI Talent Tracker*, MacroPolo (last accessed Dec. 29, 2020), <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>; Jeffrey Ding, et al., *MERICs Web Seminar: China as an AI Superpower? Quantifying China’s AI Progress Against the US and Europe*, MERICs (July 1, 2020), <https://merics.org/en/video/merics-web-seminar-china-ai-superpower-quantifying-chinas-ai-progress-against-us-and-europe>.

³ Audrey Cher, ‘Superpower Marathon’: U.S. May Lead China in Tech Right Now—But Beijing Has the Strength to Catch Up, CNBC (May 17, 2020), <https://www.cnbc.com/2020/05/18/us-china-tech-race-beijing-has-strength-to-catch-up-with-us-lead.html>; Graham Allison & Eric Schmidt, *Is China Beating the U.S. to AI Supremacy?*, Belfer Center for Science and International Affairs (Aug. 2020), <https://www.belfercenter.org/publication/china-beating-us-ai-supremacy>; Will Knight, *China May Overtake the US with the Best AI Research in Just Two Years*, MIT Technology Review (March 13, 2019), <https://www.technologyreview.com/2019/03/13/136642/china-may-overtake-the-us-with-the-best-ai-research-in-just-two-years/>.

⁴ Karen Hao, *Three Charts Show How China’s AI Industry Is Propped Up by Three Companies*, MIT Technology Review (Jan. 22, 2019), <https://www.technologyreview.com/2019/01/22/137760/the-future-of-chinas-ai-industry-is-in-the-hands-of-just-three-companies/>.

⁵ James Kynge & Nian Liu, *From AI to Facial Recognition: How China Is Setting the Rules in New Tech*, Financial Times (Oct. 7, 2020), <https://www.ft.com/content/188d86df-6e82-47eb-a134-2e1e45c777b6>.

⁶ Yusho Chao, *Chinese Venture Capitalists Take a Shine to Startups Again*, Nikkei (Sept. 13, 2020), <https://asia.nikkei.com/Business/Finance/Chinese-venture-capitalists-take-a-shine-to-startups-again>. See also, *Visualizing Chinese Tech Giants Billion-dollar Acquisitions*, CB Insights (May 28, 2020), <https://www.cbinsights.com/research/bat-billion-dollar-acquisitions-infographic/>.

⁷ See, e.g., *A Chinese Tech Giant Is Setting Up an A.I. Research Lab on Amazon’s Home Turf*, CNBC (May 2, 2017), <https://www.cnbc.com/2017/05/02/tencent-ai-research-lab-seattle.html>.

⁸ See, e.g., Saheli Roy Choudhury, *Alibaba Sets Up Joint A.I. Research Center Outside China to Focus on AI*, CNBC (Feb. 28, 2018), <https://www.cnbc.com/2018/02/28/alibaba-sets-up-joint-a-i-research-lab-in-singapore.html>.

⁹ In 2019, China had the largest number of submitted and accepted papers to the Association for the Advancement of AI (AAAI), one of the longest-running AI conferences. See *Artificial Intelligence Index: 2019 Annual Report*, Stanford Institute for Human-Centered AI at 41 (2019), https://hai.stanford.edu/sites/default/files/ai_index_2019_report.pdf. The Allen Institute for AI also predicts that China is poised to overtake the U.S. in the share of top-cited, breakthrough papers in AI by 2025. See Field Cady & Oren Etzioni, *China May Overtake US in AI Research*, Allen Institute for Artificial Intelligence (March 13, 2019), <https://medium.com/ai2-blog/china-to-overtake-us-in-ai-research-8b6b1fe30595>.

¹⁰ For a selection of such strategic documents, see *AI Policy—China*, Future of Life Institute (last accessed Dec. 30, 2020), <https://futureoflife.org/ai-policy-china/>; Graham Webster, et al., *Full Translation: China’s ‘New Generation Artificial Intelligence Development Plan,’* New America (Aug. 1, 2017), <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/> (translating China’s State Council Notice on the Issuance of the New Generation Artificial Intelligence Development Plan, dated July 20, 2017).

¹¹ Gregory C. Allen, *Understanding China’s AI Strategy*, Center for a New American Security (Feb. 6, 2019), <https://www.cnas.org/publications/reports/understanding-chinas-ai-strategy>.

Chapter 9 - Endnotes

¹² Ashwin Acharya & Zachary Arnold, *Chinese Public AI R&D Spending: Provisional Findings*, Center for Security and Emerging Technology (Dec. 2019), <https://cset.georgetown.edu/wp-content/uploads/Chinese-Public-AI-RD-Spending-Provisional-Findings-1.pdf>; see also Emily Weinstein, Mapping China's Sprawling Efforts to Recruit Scientists, *Defense One* (Nov. 30, 2020), <https://www.defenseone.com/ideas/2020/11/mapping-chinas-sprawling-efforts-recruit-scientists/170373/>; David Cyranoski, *China Joins the Battle for AI Talent*, *Nature* (Jan. 17, 2018), <https://www.nature.com/articles/d41586-018-00604-6>.

¹³ *Id.*

¹⁴ U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion: China's Pursuit of Artificial Intelligence, New Materials, and New Energy* at 46, 115-116 (June 7, 2019), <https://www.uscc.gov/sites/default/files/2019-10/June%207%202019%20Hearing%20Transcript.pdf>.

¹⁵ U.S.-China Economic and Security Review Commission, *Hearing on Technology, Trade, and Military-Civil Fusion: China's Pursuit of Artificial Intelligence, New Materials, and New Energy* (June 7, 2019), <https://www.uscc.gov/sites/default/files/2019-10/June%207%202019%20Hearing%20Transcript.pdf>.

¹⁶ *Id.*

¹⁷ As Eric Schmidt noted in *Building a New Technological Relationship and Rivalry*. See Hal Brands & Francis J. Gavin, *COVID-19 and World Order: The Future of Conflict, Competition, and Cooperation*, Johns Hopkins University Press at 406-418 (Aug. 31, 2020), <https://muse.jhu.edu/chapter/2696578>.

¹⁸ Ishan Banerjee & Matt Sheehan, *America's Got AI Talent: US' Big Lead in AI Research Is Built on Importing Researchers*, *MacroPolo* (June 9, 2020), <https://macropolo.org/americas-got-ai-talent-us-big-lead-in-ai-research-is-built-on-importing-researchers/?rp=m>.

¹⁹ *U.S. Dependence on China's Rare Earth: Trade War Vulnerability*, *Reuters* (June 27, 2019), <https://www.reuters.com/article/us-usa-trade-china-rareearth-explainer/u-s-dependence-on-chinas-rare-earth-trade-war-vulnerability-idUSKCN1TS3AQ>.

²⁰ *Member Survey*, US-China Business Council (Aug. 2019), https://www.uschina.org/sites/default/files/member_survey_2019_-_en_0.pdf.

²¹ For instance, adjusted for inflation, the Manhattan Project cost an estimated \$27 billion and the Apollo program totaled roughly \$121 billion. Deborah Stine, *The Manhattan Project, the Apollo Program, and Federal Energy Technology R&D Programs: A Comparative Analysis*, Congressional Research Service (June 30, 2009), <https://fas.org/sgp/crs/misc/RL34645.pdf> (conversion into 2020 dollars was calculated using the U.S. Bureau of Labor Statistics' CPI Inflation Calculator, available at https://www.bls.gov/data/inflation_calculator.htm).

²² In 2018, U.S. federal R&D funding amounted to 0.7% of GDP, down from its peak at above 2% in the 1970s. See James Manyika & William H. McRaven, *Innovation and National Security: Keeping Our Edge*, Council on Foreign Relations (Sept. 2019), <https://www.cfr.org/report/keeping-our-edge/recommendations/>.

²³ Zolan Kanno-Youngs & Miriam Jordan, *Trump Moves to Tighten Visa Access for High-Skilled Foreign Workers*, New York Times (Oct. 6, 2020), <https://www.nytimes.com/2020/10/06/us/politics/h1b-visas-foreign-workers-trump.html>.

²⁴ Moriah Balingit & Andrew Van Dam, *U.S. Students Continue to Lag Behind Peers in East Asia and Europe in Reading, Math and Science, Exams Show*, Washington Post (Dec. 3, 2019), https://www.washingtonpost.com/local/education/us-students-continue-to-lag-behind-peers-in-east-asia-and-europe-in-reading-math-and-science-exams-show/2019/12/02/e9e3b37c-153d-11ea-9110-3b34ce1d92b1_story.html.

²⁵ Alina Polyakova & Chris Meserole, *Exporting Digital Authoritarianism*, Brookings (Aug. 2019), <https://www.brookings.edu/research/exporting-digital-authoritarianism/>.

²⁶ The National Security Council has a statutory mandate to "advise the President with respect to the integration of domestic, foreign, and military policies relating to the national security so as to enable the Armed Forces and the other departments and agencies of the United States Government to cooperate more effectively in matters involving the national security." 50 U.S.C. § 3021(b)(1).

²⁷ See Pub. L. 94-282, National Science and Technology Policy, Organization, and Priorities Act of 1976, 90 Stat. 459 (1976), https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/ostp_organic_statute.pdf.

²⁸ The function of the NSTC under the supervision of the Director of OSTP is: "(1) to coordinate the science and technology policy-making process; (2) to ensure science and technology policy decisions and programs are consistent with the President's stated goals; (3) to help integrate the President's science and technology policy agenda across the Federal Government; (4) to ensure science and technology are considered in development and implementation of Federal policies and programs; and (5) to further international cooperation in science and technology. The Assistant may take such actions, including drafting a Charter, as may be necessary or appropriate to implement such functions." William J. Clinton, *Executive Order 12881: Establishment of the National Science and Technology Council* (Nov. 23, 1993), <https://www.govinfo.gov/content/pkg/WCPD-1993-11-29/pdf/WCPD-1993-11-29-Pg2450.pdf>.

²⁹ William J. Clinton, *Executive Order 12835: Establishment of the National Economic Council* (Jan. 25, 1993), <https://www.govinfo.gov/content/pkg/WCPD-1993-02-01/pdf/WCPD-1993-02-01-Pg95.pdf>.

³⁰ William J. Clinton, *Executive Order 12859: Establishment of the Domestic Policy Council* (Aug. 16, 1993), <https://www.archives.gov/files/federal-register/executive-orders/pdf/12859.pdf>.

Chapter 10: The Talent Competition

Winning the AI Talent Competition



**NDEA II With
a Focus on
Digital Skills**



**Attract and
Retain the
World's Brightest**

The United States is in a global competition for scarce AI talent.¹ The Commission is very concerned with current talent trends. The number of domestic-born students participating in AI doctorate programs has not increased since 1990, and competition for international students has accelerated, endangering the United States' ability to retain international students.² For the first time in our lifetime, the United States risks losing the competition for talent on the scientific frontiers. Cultivating more potential talent at home and recruiting and retaining more existing talent from foreign countries are the only two options to sustain the U.S. lead.



“For the first time in our lifetime, the United States risks losing the competition for talent on the scientific frontiers.”

Competitors and allies recognize the importance of implementing AI talent strategies. Between 2000 and 2014, China's university system increased its number of science, technology, engineering, and mathematics (STEM) graduates by 360%, producing 1.7 million in 2014 alone.³ The number of STEM graduates in the United States' university system rose by 54% during the same time period, and many were international students.⁴ China's researchers now represent roughly 29% of top-tier deep learning talent in the world.⁵ China and other states have also taken steps to attract international talent with flexible immigration policies and incentives for tech talent.⁶

The United States needs to invest in all AI talent pipelines in order to remain at the forefront of AI now and into the future. A passive strategy will not work in the face of the AI talent competition.

To achieve dominance in AI, the U.S. needs to train four archetypes to propel AI in America: researchers, implementers, end users, and informed consumers.



Researchers

AI research engineers will focus on R&D of technologies that enable and advance semi- and fully-autonomous systems. They serve as algorithm experts with up-to-date knowledge of modern AI research and may be involved in the inception of ideas and drive the development cycles from research to testing of prototypes for a major project or component of a major project.



Implementers

They will be responsible for data cleaning, feature extraction and selection, and analysis; model training and tuning; partnerships with domain knowledge experts and end users; and the discovery of local opportunities for exploitation. Developers require less training and education than AI experts, and will have training, education, and/or experience that is roughly equivalent to an associate or bachelor's degree; and that includes relevant ethics and bias mitigation in data processing and model training.



End Users

They will have their daily business augmented/enabled by AI. Use of AI will strongly resemble the use of currently available software in that it will require some system-specific training, but, with the exception of some positions that manage data, little to no AI specific expertise.



Informed Consumers

This group of people needs the ability to make better consumer choices when purchasing technology and understand the importance of their actions in the market.

The Promise and Limits of Expanding STEM.

Investments in STEM education are a necessary part of increasing American national power and improving national security. The United States ranks well overall on international measures of talent because of our ability to attract international talent, in spite of our uneven kindergarten to 12th grade (K-12) education system.⁷ It is critical that the United States invest significantly in STEM education as an engine to drive the growth of AI talent in America. Investments in STEM education alone, however, will not be enough for the United States to win the international competition for AI and STEM talent. China is producing large numbers of computer scientists, engineers, and other STEM graduates.⁸ For the foreseeable future, the United States' STEM education system does not have the capacity nor the quality to produce sufficient STEM or AI talent to supply the United States' markets or national security enterprise.⁹ To compete, the United States must reform its education system to produce both a higher quality and quantity of graduates.

Recommendation


Pass a National Defense Education Act II. Motivated by fear that America had fallen behind in education and innovation after the Soviets launched Sputnik in 1957, Congress passed the National Defense Education Act (NDEA) the following year. The NDEA promoted the importance of science, mathematics, and foreign languages for students, authorizing more than \$1 billion toward decreasing student loans, funding for education at all levels, and funding for graduate fellowships. Many students were able to attend college because of this legislation. In 1960, 3.6 million students attended college; by 1970 it was 7.5 million.¹⁰ This act helped America win the Space Race, helped power the microelectronics industry, accelerated the U.S. capacity to innovate, and, ultimately, played an important role in America's victory in the Cold War.

The Commission believes the time is right for a second NDEA, one that mirrors the first legislation, but with important distinctions. NDEA II would focus on funding students acquiring digital skills, like mathematics, computer science, information science, data science, and statistics. NDEA II should include K-12 education and reskilling programs that address deficiencies across the spectrum of the American educational system, purposefully targeting under-resourced school districts. The Commission also recommends investments in university-level STEM programs with 25,000 undergraduate, 5,000 graduate, and 500 PhD-level scholarships. Undergraduate scholarships should include credit hours at community colleges to ensure more Americans have access to affordable STEM education. Ultimately, the goal of NDEA II is to widen the digital talent pool by incentivizing programs for underrepresented Americans.

“The Commission believes the time is right for a second NDEA ...”

Recommendation

Strengthen AI talent through immigration. Immigration reform is a national security imperative. Nations that can successfully attract and retain highly skilled individuals gain strategic and economic advantages over competitors. Human capital advantages are particularly significant in the field of AI, where demand for talent far exceeds supply.¹¹ Highly skilled immigrants accelerate American innovation, improve entrepreneurship, and create jobs.¹² The United States benefits far more from the immigration of highly skilled foreign workers than other countries. In 2013, the United States had 15 times as many immigrant inventors as there were American inventors living abroad.¹³ By contrast, Canada, Germany, and the U.K. all maintain a net negative inventor immigration rate.¹⁴ Compared with other U.S. advantages in the AI competition—such as financial resources or hardware capacity—this immigration advantage is harder for other countries to replicate.



“Nations that can successfully attract and retain highly skilled individuals gain strategic and economic advantages over competitors.”

Unfortunately, international students in the United States are increasingly choosing to study in other countries or return home.¹⁵ One reason is the growing backlog of green card petitions.¹⁶ Indian immigrants face a particularly long wait. Many will spend decades on constrictive work visas waiting to receive their green cards, hindering both the technology sector's ability to recruit talent and Indian immigrants' quality of life. At the same time, other countries are increasing their efforts to attract and retain AI talent, including immigrants in Silicon Valley.¹⁷



“Restrictions harm U.S. innovation and economic growth and only help our competitors by enabling their human capital to grow. They also incentivize U.S. technology companies to move to where talent resides, whether right across our borders or overseas.”

While immigration benefits the United States, policymakers must also bear in mind the threat of unwanted technology transfer. However, restricting immigration is far too blunt a tool to solve this problem.¹⁸ Restrictions harm U.S. innovation and economic growth and only help our competitors by enabling their human capital to grow. They also incentivize U.S. technology companies to move to where talent resides, whether right across our borders or overseas.¹⁹ Technology transfer will only get worse if significant components of the U.S. technology sector move their research and development to China or other countries that are more vulnerable than the United States to technology transfer efforts.²⁰ A more effective strategic approach would pair actions to improve the United States' ability to attract top global talent with targeted efforts to combat technology transfer vectors. NSCAI addresses technology transfer in detail in Chapter 14 of this report. Changes to immigration policies should be paired with those recommendations.

Immigration policy can also slow China's progress. China's government takes the threat of brain drain seriously, noting that the United States' ability to attract and retain China's talent is an obstacle to the Chinese Communist Party's (CCP) ambitions.²¹ Increasing China's brain drain will create a dilemma for the CCP—which will be forced to choose between losing even more human capital, further slowing their economic growth and threatening their advancement in AI, or denying Chinese nationals opportunities to study and work in the United States. At the same time, the United States should be cautious about potential adverse effects on talent pools in partner nations.

“Increasing China’s brain drain will create a dilemma for the CCP ...”

Recommendation

Broaden the scope of “extraordinary” talent to make the O-1 visa more accessible and emphasize AI talent. The O-1 temporary worker visa is for people with extraordinary ability or achievement. Currently adjudicators determine an applicant's eligibility through a subjective assessment. For the sciences and technology, this aligns largely with academic criteria such as publications in major outlets and is not well suited for people who excel in industry.

Recommendation

Implement and advertise the international entrepreneur rule. The International Entrepreneur Rule (IER) allows U.S. Citizenship and Immigration Services (USCIS) to grant a period of authorized stay to international entrepreneurs who demonstrate that “their stay in the United States would provide a significant public benefit through their business venture.”²² An executive action could announce the administration's intention to use the IER to boost immigrant entrepreneurship, job creation for Americans, and economic growth. USCIS

could also be directed to announce that it will give priority to entrepreneurs active in high-priority STEM fields such as AI, or in fields that use AI for other applications, such as agriculture. Entrepreneurs' ability to attract investors should be used as a screening criterion for entrepreneurs.

Expand and clarify job portability for highly skilled workers. The criteria for workers with H-1B, O-1, and other temporary work visas to obtain open market work permits for a one-year renewable period are too limited and ambiguous. Changes should clarify when highly skilled, nonimmigrant workers are permitted to change jobs or employers, increase job flexibility when an employer either withdraws their petition or goes out of business, and increase flexibility for H-1B workers seeking other H-1B employment.

Recommendation

Recapture green cards lost to bureaucratic error. Federal agencies generally issue fewer green cards than they are allowed. As of 2009, the federal government had failed to issue more than 326,000 green cards based on cumulative bureaucratic error.²³ The Departments of State and Homeland Security (DHS) should publish an up-to-date report on the number of green cards lost due to bureaucratic error. Using available authorities, both should grant lost green cards to applicants waiting in line. Congress should support the Departments of State and Homeland Security by passing legislation to recapture lost green cards.²⁴

Recommendation

Grant green cards to students graduating with STEM PhDs from accredited American universities. Congress should amend the Immigration and Nationality Act²⁵ to grant lawful permanent residence to any vetted (not posing a national security risk) foreign national who graduates from an accredited United States institution of higher education with a doctoral degree in a STEM-related field in a residential or mixed residential and distance program and has a job offer in a field related to science, technology, engineering, or mathematics. They should not be counted toward permanent residency caps.

Recommendation

Double the number of employment-based green cards. Under the current system, employment-based green cards are unduly scarce: 140,000 per year, fewer than half of which go to the principal worker.²⁶ This leaves many highly skilled workers unable to gain permanent residency and unable to transfer jobs or negotiate with employers as effectively as domestic workers. This decreases the appeal of joining the American workforce. To reduce the backlog of highly skilled workers, the United States should double the number of employment-based green cards, with an emphasis on permanent residency for STEM and AI-related fields.

Recommendation

Create an entrepreneur visa. International doctoral students are more likely than their native peers to want to found a company or become an employee at a startup, but they are less likely to pursue those paths.²⁷ One reason is the constraints of the H-1B visa system.²⁸ Similarly, immigrant entrepreneurs without the capital to use the EB-5 route to permanent residency are forced to use other visas that are designed for academics and workers in existing companies, not entrepreneurs.²⁹ All of these issues make the United States

Recommendation

less attractive for international talent, and, perhaps as important, reduce the ability of startups and other small companies—the main source of new jobs for Americans—to hire highly skilled immigrants, who have been shown to improve the odds that the business will succeed. Congress should create an entrepreneur visa for those who would provide a “significant public benefit” to the United States if allowed to stay in the country for a limited trial period to grow their companies.³⁰ This visa should serve as an alternative to employee-sponsored, investor, or student visas and should instead target promising potential founders.

Recommendation

Create an emerging and disruptive technology visa. The National Science Foundation (NSF) should identify critical emerging technologies every three years. DHS would then allow students, researchers, entrepreneurs, and technologists in applicable fields to apply for emerging and disruptive technology visas. This would provide much-needed talent R&D and strengthen our economy.³¹

Chapter 10 - Endnotes

¹ Estimates on the gap of talent necessary to fill AI slots vary greatly, but it is agreed upon that the gap in talent currently is and will continue to be significant as nations compete for scarce resources. See Remco Zwetsloot, et al., *Strengthening the U.S. AI Workforce: A Policy and Research Agenda*, Center for Security and Emerging Technology at 2 (Sept. 2019), <https://cset.georgetown.edu/wp-content/uploads/CSET-Strengthening-the-U.S.-AI-Workforce.pdf> (“The Research Institute at Tencent, a major Chinese technology company, asserts there are roughly 300,000 AI researchers and practitioners worldwide, with market demand for millions of roles. Element AI, a leading Canadian AI company, estimated in 2018 that there are roughly 22,000 PhD-educated researchers globally who are able to work on AI research, with only about 25 percent of those ‘well-versed enough in the technology to work with teams to take it from research to application.’ AI firm Diffbot estimates that there are over 700,000 people skilled in machine learning worldwide.”).

² Remco Zwetsloot, et al., *Keeping Top AI Talent in the United States*, Center for Security and Emerging Technology at iii-vi (Dec. 2019), <https://cset.georgetown.edu/wp-content/uploads/Keeping-Top-AI-Talent-in-the-United-States.pdf>.

³ *The Rise of China in Science and Engineering*, NSF National Science Board (2018), <https://www.nsf.gov/nsb/sei/one-pagers/China-2018.pdf> (China also passed the United States in the global share of peer-reviewed S&E articles).

⁴ *Science & Engineering Indicators 2018*, NSF National Science Board (2018), <https://www.nsf.gov/statistics/2018/nsb20181/assets/561/higher-education-in-science-and-engineering.pdf>.

⁵ For these purposes “top tier” talent was defined by accepted papers at the prestigious AI deep learning conference Neural Information Processing Systems in 2019. This reflected approximately the top 20% of researchers in the field. *The Global AI Talent Tracker*, MacroPolo (last accessed Dec. 28, 2020), <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>. China has placed a strong emphasis on deep learning, just one of the important components of AI.

⁶ For example, China’s Thousand Talents Plan is part of a state-organized blueprint to be a global leader in science and technology by 2050. Staff Report, *Threats to the U.S. Research Enterprise: China’s Talent Recruitment Plans*, U.S. Senate Permanent Subcommittee on Investigations at 14 (Nov. 2019), <https://www.hsgac.senate.gov/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China%20Talent%20Recruitment%20Plans.pdf>.

⁷ *The Global AI Talent Tracker*, MacroPolo (last accessed Dec. 28, 2020), <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>. See also Gordon Hanson & Matthew Slaughter, *High-Skilled Immigration and the Rise of STEM Occupations in U.S. Employment*, National Bureau of Economic Research at 1 (Sept. 2016), https://www.nber.org/system/files/working_papers/w22623/w22623.pdf.

⁸ *The Rise of China in Science and Engineering*, NSF National Science Board (2018), <https://www.nsf.gov/nsb/sei/one-pagers/China-2018.pdf>.

⁹ As noted in Chapter 6 of this report, there were 433,116 open computer science jobs in the United States in 2019, while only 71,226 new computer scientists graduated from American universities in 2019. Code.org (last accessed Jan. 11, 2021), <https://code.org/promote>. See also Oren Etzioni, *What Trump’s Executive Order on AI Is Missing: America Needs a Special Visa Program Aimed at Attracting More AI Experts and Specialists*, *Wired* (Feb. 13, 2019), <https://www.wired.com/story/what-trumps-executive-order-on-ai-is-missing/>.

¹⁰ *Sputnik Spurs Passage of the National Defense Education Act*, U.S. Senate (last accessed Dec. 28, 2020), https://www.senate.gov/artandhistory/history/minute/Sputnik_Spurs_Passage_of_National_Defense_Education_Act.htm#:~:text=The%20National%20Defense%20Education%20Act%20of%201958%20became%20one%20of.and%20private%20colleges%20and%20universities.

¹¹ According to one report, job listings for AI on one popular job website “increased more than five-fold between 2015 and 2017 and demand for ‘deep learning’ skills increased by a factor of more than 30,” while the number of people looking for jobs in the field grew much more slowly. This mismatch is slowing the adoption of AI. Most firms report that skills gaps are one of the top obstacles preventing them from adopting AI. Remco Zwetsloot, et al., *Strengthening the U.S. AI Workforce: A Policy and Research Agenda*, Center for Security and Emerging Technology at 1 (Sept. 2019), <https://cset.georgetown.edu/wp-content/uploads/CSET-Strengthening-the-U.S.-AI-Workforce.pdf>.

Chapter 10 - Endnotes

- ¹² William S. Kerr, *High-Skilled Immigration, Innovation, and Entrepreneurship: Empirical Approaches and Evidence*, National Bureau of Economic Research at 7-8 (Aug. 2013), <https://www.nber.org/papers/w19377>; Gordon Hanson & Matthew Slaughter, *Strengthening the U.S. AI Workforce, High-Skilled Immigration and the Rise of STEM Occupations in U.S. Employment*, National Bureau of Economic Research at 23 (Sept. 2016), https://www.nber.org/system/files/working_papers/w22623/w22623.pdf; Remco Zwetsloot, et al., *Strengthening the U.S. AI Workforce: A Policy and Research Agenda*, Center for Security and Emerging Technology at 5 (Sept. 2019), <https://cset.georgetown.edu/wp-content/uploads/CSET-Strengthening-the-U.S.-AI-Workforce.pdf>.
- ¹³ Carsten Fink, *What Leads Inventors to Migrate?*, World Economic Forum (July 17, 2013), <https://www.weforum.org/agenda/2013/07/what-leads-inventors-to-migrate/>.
- ¹⁴ Ernest Miguelez & Carsten Fink, *Measuring the International Mobility of Inventors: A New Database*, World Intellectual Property Organization at 16 (May 2013), https://www.wipo.int/edocs/pubdocs/en/wipo_pub_econstat_wp_8.pdf.
- ¹⁵ According to the Center for Security and Emerging Technology, in 2016, 14% of international students declined offers to study at U.S. universities to study at home, and 19% decided to study in another country. In 2018, these numbers rose, with 39% staying at home and 59% studying in another country. Remco Zwetsloot, et al., *Keeping Top AI Talent in the United States: Findings and Policy Options for International Graduate Student Retention*, Center for Security and Emerging Technology at 26 (Dec. 2019), <https://cset.georgetown.edu/wp-content/uploads/Keeping-Top-AI-Talent-in-the-United-States.pdf>.
- ¹⁶ Shulamit Kahn & Megan MacGarvie, *The Impact of Permanent Residency Delays for STEM PhDs: Who Leaves and Why*, Research Policy (Nov. 2020), <https://www.sciencedirect.com/science/article/abs/pii/S0048733319301982>.
- ¹⁷ Tina Huang & Zachary Arnold, *Immigration Policy and the Global Competition for AI Talent*, Center for Security and Emerging Technology at 8 (June 2020), <https://cset.georgetown.edu/research/immigration-policy-and-the-global-competition-for-ai-talent/>.
- ¹⁸ Zachary Arnold, et al., *Immigration Policy and the U.S. AI Sector: A Preliminary Assessment*, Center for Security and Emerging Technology at 22 (Sept. 2019), <https://cset.georgetown.edu/research/immigration-policy-and-the-u-s-ai-sector/>.
- ¹⁹ Remco Zwetsloot, et al., *Strengthening the U.S. AI Workforce: A Policy and Research Agenda*, Center for Security and Emerging Technology at 5 (Sept. 2019), <https://cset.georgetown.edu/wp-content/uploads/CSET-Strengthening-the-U.S.-AI-Workforce.pdf>.
- ²⁰ China is the world's largest single source of AI talent. Leading U.S. technology companies such as Google and Microsoft have established cutting-edge research centers in China, in part to access that talent. This increases China's AI R&D capacity and potential for technology transfer, and, if the companies remain American, it reduces the American Intelligence Community's (IC) legal authorization to collect information about Chinese technology development. See *The Global AI Talent Tracker*, MacroPolo (last accessed Jan 17, 2020), <https://macropolo.org/digital-projects/the-global-ai-talent-tracker/>; Roxanne Heston & Remco Zwetsloot, *Mapping U.S. Multinationals' Global AI R&D Activity*, Center for Security and Emerging Technology at 20 (Dec. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Mapping-U.S.-Multinationals-Global-AI-RD-Activity-1.pdf>.

²¹ Remco Zwetsloot, *US-China STEM Talent "Decoupling": Background, Policy, and Impact*, Johns Hopkins Applied Physics Laboratory at 19 (2020), <https://www.jhuapl.edu/assessing-us-china-technology-connections/dist/407b0211ec49299608551326041488d4.pdf> ("[T]he head of the [Chinese Communist Party's (CCP)] Central Talent Work Coordination Small Group ... complained that 'the number of top talents lost in China ranks first in the world.'"); see also Joy Dantong Ma, *China's AI Talent Base Is Growing, and Then Leaving*, MacroPolo (July 30, 2019), <https://macropolo.org/chinas-ai-talent-base-is-growing-and-then-leaving/?rp=m> (noting that of the 2,800 Chinese NeurIPS participants between 2009 and 2018, about three-quarters of them were currently working outside of China).

²² *International Entrepreneur Parole*, USCIS (May 25, 2018), <https://www.uscis.gov/humanitarian/humanitarian-parole/international-entrepreneur-parole>. There is currently no visa category well-suited to entrepreneurship in immigration statute. The IER, which relies on parole authority, was initiated after legislative avenues were exhausted. Legislative fixes would be preferable, but have so far they have proven politically infeasible.

²³ A 2010 report to Congress indicated that some 242,000 unused family-based green cards were ultimately applied to the employment-based backlog, while Congress recaptured some 180,000 green cards via special legislation, leaving more than 326,000 green card numbers wasted. *Citizenship and Immigration Services Ombudsman: Annual Report 2010*, U.S. Department of Homeland Security (June 30, 2010), https://www.dhs.gov/xlibrary/assets/cisomb_2010_annual_report_to_congress.pdf. The number today is likely higher, but DHS has not published updated statistics.

²⁴ Prior examples of Congressional action include provisions in the American Competitiveness in the 21st Century Act of 2000 and the REAL ID Act of 2005. See Pub. L. 106-313, 114 Stat. 1251, 1254 (2000) and Pub. L. No. 109-013, 119 Stat. 231, 322 (2005).

²⁵ Specifically, 8 U.S.C. § 1151(b)(1).

²⁶ William Kandel, *The Employment-Based Immigrant Backlog*, Congressional Research Service at 4-5 (March 26, 2020), <https://fas.org/sgp/crs/homesecc/R46291.pdf>.

²⁷ Michael Roach, et al., *Are Foreign STEM PhDs More Entrepreneurial? Entrepreneurial Characteristics, Preferences and Employment Outcomes of Native and Foreign Science & Engineering PhD Students*, National Bureau of Economic Research at 1 (Sept. 2019), https://www.nber.org/system/files/working_papers/w26225/w26225.pdf.

²⁸ *Id.* at 12.

²⁹ EB-5 visas require a minimum \$900,000 investment in a business in the United States. William R. Kerr, *Global Talent and U.S. Immigration Policy: Working Paper 20-107*, Harvard Business School at 14 (2020), https://www.hbs.edu/faculty/Publication%20Files/20-107_0967f1ab-1d23-4d54-b5a1-c884234d9b31.pdf.

³⁰ 83 Fed. Reg. 24415, *Removal of International Entrepreneur Parole Program*, U.S. Department of Homeland Security (May 29, 2018), <https://www.federalregister.gov/documents/2018/05/29/2018-11348/removal-of-international-entrepreneur-parole-program>.

³¹ Oren Etzioni, *What Trump's Executive Order on AI Is Missing: America Needs a Special Visa Program Aimed at Attracting More AI Experts and Specialists*, *Wired* (Feb. 13, 2019), <https://www.wired.com/story/what-trumps-executive-order-on-ai-is-missing/>.

Chapter 11: Accelerating AI Innovation



Leadership in AI Innovation

Scale and
Coordinate Federal
AI R&D Funding



Expand Access
Through a National AI
R&D Infrastructure



Strengthen
Public-Private
Partnerships



Tackle Humanity's
Biggest Challenges



To remain the world's leader in artificial intelligence (AI), the U.S. government must renew its commitment to investing in America's national strength: innovation. This will require making substantial new investments in AI R&D and establishing a national AI research infrastructure that democratizes access to the resources that fuel AI. Members of Congress must come to terms with the reality that tens of billions of dollars will be needed over the next several years. The return on these investments will transform America's economy, society, and national security.

A lack of national urgency is dangerous at a time when underlying weaknesses have emerged in our AI ecosystem that impair innovation and when viewed against the backdrop of China's state-directed AI progress. The development of AI in the United States is concentrated in fewer organizations in fewer geographic regions pursuing fewer research pathways. Commercial agendas are dictating the future of AI and concentrating heavily in one discipline: machine learning (ML).¹ Despite promising moves, government funding has lagged behind the transformative potential of the field, limiting its ability to shape research toward the public good and support progress across a range of AI disciplines.² As a result, the AI innovation environment rests on a narrowing foundation.



“A lack of national urgency is dangerous at a time when underlying weaknesses have emerged in our AI ecosystem that impair innovation, and when viewed against the backdrop of China’s state-directed AI progress.”

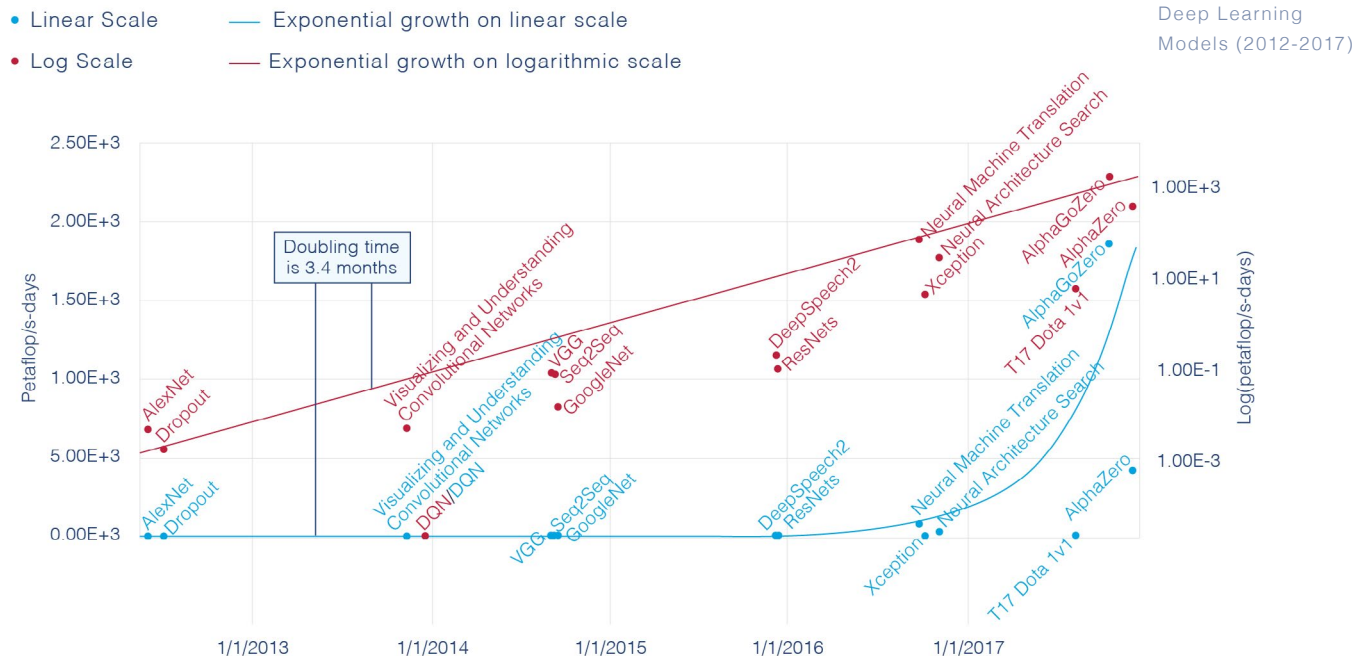
These trends toward consolidation come as a result of resources. Declining per-unit costs of cloud-based computing and availability of open-source platforms have lowered barriers to access for core ML. However, those same conditions have enabled pursuit of sophisticated models that require extensive training data, often held in privately controlled data sets or knowledge graphs, enormous computing power, and substantial hardware and software engineering.³ These prerequisites now define the cutting edge of AI research and effectively limit the number of American researchers able to contribute to the field and tackle the hard challenges that could unlock new frontiers of AI.

Ingenuity, Not Access, Should Be the Key to AI Innovation in America.

The consolidation of the AI industry threatens U.S. technological competitiveness in five ways:

- **Brain Drain.** Brain drain from academic institutions to the private sector threatens to hollow out the foundations of the United States' advantage in basic AI research: its universities.⁴ Federal funding that has not kept pace with the growth of the field has led to low grant application success rates and amplified time spent on the bureaucracy of pursuing and completing proposals.⁵ However, academic experts and their students are not just lured to big tech by the promise of less bureaucracy and higher financial incentives. Increasingly, the private sector is the best place to conduct cutting-edge research with access to the best computing and data resources. The result is the weakening of the teaching base for the next generation of AI leaders in industry and academia and the narrowing of the overall AI research agenda.⁶

Compute Required to Train Largest Deep Learning Models (2012-2017)



A "petaflop/s-day" is a measure of compute that consists of performing 10¹⁵ floating point operations per second for one day.

GPT-3 and more recent models are not represented here because they are too large to fit on this scale.

- **Diversity.** The growing divide between “haves” and “have nots” in AI will exacerbate the well-documented lack of diversity in the field,⁷ limiting the field’s collective ability to build equitable, inclusive systems.
- **Research Focus.** American technology firms are accountable to their shareholders and will logically not invest in areas of national security importance or make uncertain bets on fundamental research that does not hold commercial or economic benefit for the company.⁸ While return-focused investments can lead to applications that contribute to the public good or benefit government work, there are gaps. ML and the underlying algorithms were in exactly this position two decades ago—seemingly without commercial promise—only to be sustained by federal research dollars until computing power and an overabundance of data transformed the discipline.⁹ A recent study found that 82% of the algorithms in use today originated from federally funded non-profits and universities, compared to just 18% that originated from private companies.¹⁰
- **Competition.** The rising cost of developing cutting-edge ML models and high likelihood of acquisition by leading technology companies means AI startups have narrowing paths to growth in the United States.¹¹ Lack of competition undermines the industry’s ability to innovate and be globally competitive in the research and development of AI.
- **Regional Divergence.** The clustering of technology firms in regions like Silicon Valley drives innovation by expediting knowledge sharing and sharpening domestic rivalry.¹² However, this trend has benefited some regions and demographics more than others.¹³ More than 90% of U.S. innovation sector job creation occurred in just five major coastal cities between 2005 and 2017.¹⁴ This divergence concentrates gains from technological progress in just a few regions and misses out on latent innovation potential in the rest of the country.

The federal government holds the responsibility to reverse these trends. It must step in and step up to provide strategic direction and sustained resources, as both a funder and consumer of technology.¹⁵ It must break the mold of standard scientific research funding. The outcomes of technology innovation, which generate greatest value when translated into fieldable solutions, are driven by multi-sector contributions and a culture of risk acceptance. The status quo at federal agencies and research entities is insufficient to make these big bets and propel promising technology concepts from laboratory to field.

A passive national approach that relies too heavily on the private sector to drive innovation and determine research agendas—and that presumes commercial innovation can simply “spin-in” to become government applications—will not win this strategic competition, nor will it fully capitalize on the transformative potential of AI. The United States—through government leadership in partnership with industry and academia—must increase the diversity, competitiveness, and accessibility of its AI innovation environment. That begins with a substantial infusion of new R&D dollars.

“The United States—through government leadership in partnership with industry and academia—must increase the diversity, competitiveness, and accessibility of its AI innovation environment.”


Scale and coordinate federal AI R&D funding. A bold, integrated push for long-term investments in AI R&D will foster nationwide AI innovation and drive breakthroughs. An infusion of sustained resources, guided by a comprehensive strategy and distributed through a diversity of mechanisms, will enable U.S. researchers to push the boundaries of the field by supporting a wide range of AI approaches and novel applications of AI to other fields. Specifically, the United States should:

Recommendation

- **Establish a National Technology Foundation (NTF).** A new, independent organization would complement successful existing organizations, such as the NSF and DARPA, by providing the means to more aggressively move science into engineering. The NTF would drive technology progress at a national level by focusing on generating value at intermediate levels of technical maturity, prioritizing use-inspired concepts, establishing infrastructure for experimentation and testing, and supporting commercialization of successful outcomes. This requires an organization that is structured to accept higher levels of risk and empowered to make big bets on innovative ideas and people.
- **Increase federal funding for non-defense AI R&D at compounding levels, doubling annually to reach \$32 billion per year by Fiscal Year 2026.** This would bring AI spending to a level near to federal spending on biomedical research.¹⁶ Overall, the government should spend at least 1% of GDP on R&D to reinforce a base of innovation across scientific fields.¹⁷ Additional funding should strengthen basic and applied research, expand fellowship programs, support research infrastructure, and cover several agencies, with an emphasis on:
 - o National Technology Foundation (proposed)
 - o Department of Energy
 - o National Science Foundation
 - o National Institutes of Health
 - o National Institute of Standards and Technology
 - o National Aeronautics and Space Administration

AI R&D
Investment Levels.

	FY2020	FY2021	
NSF	\$518.3M	\$831.2M	32
NIH	\$193.9M	\$176.8M	30
DOE	\$171.8M	\$174.4M	28
USDA	\$54.9M	\$129.6M	26
DHS	\$50.4M	\$31.3M	24
FDA	\$39.0M	\$38.0M	22
NASA	\$28.5M	\$28.8M	20
NIST	\$27.6M	\$52.7M	18
DOT	\$17.1M	\$16.3M	16
VA	\$14.1M	\$14.1M	14
DOI	\$5.9M	\$4.2M	12
NIJ	\$3.0M	\$3.0M	10
NOAA	\$1.6M	\$1.6M	8
Treasury	\$0.6M	\$0.6M	6
Total	\$1.127B	\$1.503B	4



DHS - Department of Homeland Security
DOE - Department of Energy
DOI - Department of the Interior
DOT - Department of Transportation
FDA - Food and Drug Administration
NASA - National Aeronautics and Space Administration
NIH - National Institutes of Health
NIJ - National Institute of Justice

NIST - National Institute of Standards and Technology
NOAA - National Oceanic and Atmospheric Administration
NSF - National Science Foundation
Treasury - Treasury/Financial Crimes Enforcement Network
USDA - U.S. Department of Agriculture
VA - Department of Veterans Affairs

Source: The Networking & Information Technology Research & Development Program, Supplement to The President's FY2021 Budget, National Science & Technology Council (Aug. 14, 2020), <https://www.nitrd.gov/pubs/FY2021-NITRD-Supplement.pdf>.

- **Prioritize funding for key areas of AI R&D.** Amplified federal funding should prioritize AI R&D investments in areas critical to advance technology that will underpin future national security and economic stability, supporting areas that may not receive significant private-sector investment. Coordinated through the newly established National AI Initiative,¹⁸ investments should reflect a portfolio approach, focused on advancing basic science, solving specific challenge problems, and facilitating commercialization breakthroughs.

Priority Areas for AI Research Investment.



- **Triple the number of National AI Research Institutes.** The government should triple the current number of federally funded national AI research institutes across a range of regions and research areas.¹⁹ This would increase training and research opportunities for students and academic faculty, national lab researchers, and non-profit research organizations.
- **Invest in talent that will transform the field.** In parallel, NSF or the proposed NTF should invest in top AI researchers and interdisciplinary teams, launching grant awards that make big bets on the people and the out-of-the-box ideas that could lead to unexpected breakthroughs.

“Democratized access to compute environments, data, and testing facilities will provide researchers beyond leading industry players and elite universities the ability to pursue progress on the cutting edge of AI.”

Recommendation


Expand access to AI resources through a National AI Research Infrastructure. Democratized access to compute environments, data, and testing facilities will provide researchers beyond leading industry players and elite universities the ability to pursue progress on the cutting edge of AI. It will strengthen the foundation of American AI innovation by supporting more equitable growth of the field, expanding AI expertise across the country, and applying AI to a broader range of fields. This national infrastructure should have five main elements:

- **A National AI Research Resource (NAIRR).**²⁰ To bridge the “compute divide,”²¹ the NAIRR would provide verified researchers and students subsidized access to scalable compute resources, co-located with AI-ready government and non-government data sets, educational tools, and user support. It should be created as a public-private partnership, leveraging a federation of cloud platforms.²²
- **A set of domain-specific AI R&D test beds.** Sponsored by various federal agencies, these would provide accessible facilities, establish benchmarking standards, and build communities of discovery around AI application areas that are in the public interest.

- **Large-scale, open training data.** This should include curation, hosting, and maintenance of complex data sets; incentives to the private sector and academia to share data sets; and funding for teams of data engineers and scientists to unlock public data currently held by the government for use by the AI research community.
- **An open knowledge network.** Coordinated by the Office of Science and Technology Policy, such a resource would enable use of constructed and organized world knowledge to develop AI systems that operate effectively and efficiently.²³
- **A Multilateral AI Research Institute.** To foster collaborative R&D with researchers from key allies and partners (described further in Chapter 15 of this report).

These resources would work in complement to each other, providing a virtuous cycle of data, experimentation, testing, and knowledge-building that would fuel innovation and application of AI to a wide range of challenge problems and fields of study.

National AI
Research
Infrastructure.



Leverage both sides of the public-private partnership. U.S. leadership in technologies like AI depends upon closer public-private collaboration and a shared sense of responsibility for U.S. global competitiveness. To promote innovation and accelerate the growth of globally competitive firms in strategic emerging sectors, the government should:

Recommendation


- **Create markets for AI and other strategic technologies.** The application of AI across government agencies can save taxpayer dollars and improve the quality of public

services. Some applications can be adopted directly from the private sector, while others are unique to the government mission. By accelerating AI adoption across federal agencies, the government can drive additional commercial investment in AI applications that benefit national security and the public good.²⁴

- **Form a network of regional innovation clusters focused on strategic emerging technologies.** The government should designate regional innovation clusters focused on strategic emerging technologies like AI to foster the growth of small companies in sectors that are critical to overall U.S. competitiveness. Established through a competitive bid process, the clusters would offer participants from industry and academia tax incentives, research grants, and access to federal R&D resources.

Regional Innovation Clusters.

Cluster Strength* by County, 2017



*Cluster strength is the percent of traded employment in a region with high employment specialization. This is one of many important factors to consider when selecting locations for regional innovation clusters.


Image source: U.S. Cluster Mapping Project, Institute for Strategy and Competitiveness, Harvard Business School. Data source: U.S. Census Bureau.

The private sector should:

- **Privately fund an AI competitiveness consortium.** Private firms should establish a non-profit organization with \$1 billion in funding over five years to broaden AI research opportunities and support AI skills and education. This donation-funded organization would focus on fostering economic opportunity through resources for AI research and AI skills training. Such corporate social responsibility spending to promote AI education and entrepreneurship would help bridge the gap between digital “haves” and “have nots.”

Tackle some of humanity's biggest challenges.

Recommendation



Tackle Some of Humanity's Biggest Challenges.

By focusing on solving real human problems that impact the lives of millions of people, we can build a new raison d'être for the triangular alliance of government, academia, and industry; sustain public support for ambitious AI research; and extend America's AI innovation leadership. Examples of promising initiatives that could improve societal well-being and advance scientific frontiers include:

Enable long-term quality of life.

- AI technology that can help the elderly live independently longer, assisting in managing health and daily tasks and improving the quality of life
- This can include application of AI to biomedicine to address acute and chronic illnesses and enhance healthy aging



Revolutionize education and life long learning. AI tools that personalize education, training, and retraining at appropriate challenge levels and intuitively evaluate development to optimize standard curricula to promote individual learning success

Transform energy management. Smart infrastructure for cities that can effectively respond to surges in energy demand and emergencies (both man-made and natural disasters)



Effectively predict, model, prepare for – and respond to disasters.

- Accurate, near-real time weather, earthquake, and fire line detection and prediction of escalation to aid in emergency response and planning for optimized deployment of limited resources
- Autonomous robots for search, rescue, and cleanup in the wake of natural or man-made disaster, providing force-multiplying support to first responders and hazardous-materials professionals



Reach new frontiers in space. Autonomous AI spacecraft, habitats, and facilities capable of identifying and solving problems with or without the need for human intervention, enabling extended and more flexible space exploration



Tackle Some of Humanity's Biggest Challenge.

“By focusing on solving real human problems that impact the lives of millions of people, we can build a new raison d'être for the triangular alliance of government, academia, and industry ...”

Chapter 11 - Endnotes

¹ A 2020 analysis of arXiv papers on AI found private-sector basic AI research to be thematically narrower than the broader corpus of AI publications, focusing on deep learning and computational infrastructure to support deep learning. Furthermore, the study found that elite academic institutions that collaborate more closely with industry had a similar narrowing of thematic concentration, leading to a tilting of the U.S. AI research environment away from the diversity still preserved in other countries. Joel Klinger, et al., *A Narrowing of AI Research?*, ArXiv (Nov. 18, 2020), <https://arxiv.org/pdf/2009.10385.pdf>. Increasing specialization of hardware achieved through industry investments has further prioritized commercial use cases, making it costly to pursue approaches outside the mainstream. Sara Hooker, *The Hardware Lottery*, ArXiv (Sept. 21, 2020), <https://arxiv.org/pdf/2009.06489.pdf>.

² The Trump Administration's proposed budget for non-defense AI R&D in Fiscal Year 2021 was \$1.5 billion, a growth from the just under \$1 billion spent in Fiscal Year 2020. *The Networking & Information Technology Research & Development Program, Supplement to The President's FY2021 Budget*, National Science & Technology Council at 4, 12 (Aug. 14, 2020), <https://www.nitrd.gov/pubs/FY2021-NITRD-Supplement.pdf>. The National AI Initiative Act of 2020 included in the National Defense Authorization Act for 2021 included authorization for additional investments in AI R&D at the National Science Foundation (NSF), Department of Energy (DOE), National Institute of Standards and Technology (NIST), and the National Oceanic and Atmospheric Administration (NOAA). See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

³ OpenAI estimates that since 2012, the amount of compute used in the largest AI training runs is doubling every 3.4 months. See Dario Amodei & Danny Hernandez, *AI and Compute*, OpenAI (May 16, 2018), <https://openai.com/blog/ai-and-compute/>. Based on projections by OpenAI, at the current rate of increasing costs of model training, "in 4 years, training the largest model will cost more than launching a rocket into orbit." Yaroslav Bulatov, *Large-scale AI and Sharing of Models*, Medium (July 20, 2019), <https://yaroslavvb.medium.com/large-scale-ai-and-sharing-of-models-4622ba59ec18>. For efforts that involve robotics or real-world applications, development requires additional resources in terms of complex modeling and simulation capabilities for training algorithms as well as specialized facilities for experimentation.

⁴ A recent study found that from 2004 to 2018, 131 AI professors left universities for industry and 90 adopted a dual affiliation while maintaining part-time positions at a university. The study also documented the adverse effect that these departures had on AI startups of students from these universities. Michael Gofman & Zhao Jin, *Artificial Intelligence, Education, and Entrepreneurship*, SSRN at 2 (Oct. 26, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3449440. High salaries in the commercial sector pull researchers from academic tracks—in 2019, 57% of AI/ML PhD graduates in North America went to industry versus staying in academia for post-doc, research, or faculty appointments. Stuart Zweben & Betsy Bizot, 2019 Taulbee Survey, Computing Research Association at 11 (May 2020), <https://cra.org/wp-content/uploads/2020/05/2019-Taulbee-Survey.pdf>.

⁵ For example, NSF, which provides 85% of federal funding for computer science, funded \$188 million in core AI research in 2019 but did not have room in the budget to fund another \$178 million worth of highly rated proposals. This was an improvement from 2018, when they funded \$165 million but left \$185 million of highly rated work unfunded. Furthermore, NSF (in partnership with the Department of Agriculture) funded seven National AI Research Institutes in 2020 but was unable to fund the more than 30 that were judged worthy of supporting; NSF presentation to NSCAI (January 2020).

⁶ The time computer science faculty can spend holding concurrent appointments in industry has increased from 20% to up to 50% to 80%, which has implications on their academic responsibilities including recruitment of students and development of coursework and seminars, as well as the possible consequence of aligning graduate-student work to industry's needs over high-impact basic research. Shwetak Patel, et al., *Evolving Academia/Industry Relations in Computing Research*, Computing Community Consortium at 3 (June 2019), <https://cra.org/ccc/wp-content/uploads/sites/2/2019/06/Evolving-AcademiaIndustry-Relations-in-Computing-Research.pdf>.

⁷ The annual Taulbee study that tracks the field of computer science (CS) found that women make up 21.0% of CS bachelor degree graduates and 20.3% of CS doctoral graduates, and domestic underrepresented minorities make up 14.7% of CS bachelor degree graduates and only 3.1% of doctoral graduates. Stuart Zweben & Betsy Bizot, 2019 Taulbee Survey, Computing Research Association at 4, 5, 22 (May 2020), <https://cra.org/wp-content/uploads/2020/05/2019-Taulbee-Survey.pdf>. A trend toward narrowing participation in the field holds the potential to worsen this state. See Nur Ahmed & Muntasir Wahed, *The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, ArXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>.

⁸ See, e.g., Joel Klinger, et al., *A Narrowing of AI Research?*, ArXiv (Nov. 18, 2020), <https://arxiv.org/pdf/2009.10385.pdf>; Sara Hooker, *The Hardware Lottery*, ArXiv (Sept. 21, 2020), <https://arxiv.org/pdf/2009.06489.pdf>.

⁹ From the very outset of the field, the federal government had a hand in fostering research. The Air Force, via RAND, supported the work of Herbert Simon and Allen Newell, who in 1956 created the first successful AI computer program, the Logic Theorist. Mariana Mazzucato, *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*, Anthem Press (2013). The Defense Advanced Research Projects Agency (DARPA) (then ARPA) funded the work of Charles Rosen, who developed the first self-navigating robot, “Shakey,” in 1972. *Shakey the Robot*, DARPA (last accessed Dec. 30, 2020), <https://www.darpa.mil/about-us/timeline/shakey-the-robot>. Reinforcement learning, the approach on which many of today’s commercial applications are based, was sustained through the “AI Winter” of the 1990s by NSF’s support of Andrew Barto; NSCAI staff engagement with NSF (Aug. 8, 2019). DARPA’s 30 years of funding for research on image understanding created the foundation for autonomous driving capabilities. DARPA, *DARPA Artificial Intelligence Colloquium Opening Video*, YouTube (March 12, 2019), <https://www.youtube.com/watch?v=FTaW6ZJ9oyQ>. The PAL program run by DARPA in the mid-2000s led to the development of the first artificially intelligent assistant, which eventually became Siri. DARPA, *DARPA and AI: Visionary Pioneer and Advocate*, YouTube (Dec. 7, 2018), <https://www.youtube.com/watch?v=ri5gOjYgLn8>.

¹⁰ Neil C. Thompson, et al., *Building the Algorithm Commons: Who Discovered the Algorithms that Underpin Computing in the Modern Enterprise?*, Global Strategy Journal at 4 (2020), <https://onlinelibrary.wiley.com/doi/epdf/10.1002/gsj.1393>.

¹¹ For example, non-elite universities and AI startups have difficulty affording the cost of compute resources and data for training sophisticated ML models. Nur Ahmed & Muntasir Wahed, *The Democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, arXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>. Ninety percent of Silicon Valley AI startups were purchased by large technology companies between 2013 and 2018. Ryan Kottenstette, *Silicon Valley Companies Are Undermining the Impact of Artificial Intelligence*, TechCrunch (March 15, 2018), <https://techcrunch.com/2018/03/15/silicon-valley-companies-are-undermining-the-impact-of-artificial-intelligence/>. These same companies dominate U.S. patent lists, excluding adoption patents. AI AuYeung, *Who is Winning the AI Race?*, IPWatchdog (Feb. 1, 2020), <https://www.ipwatchdog.com/2020/02/01/winning-ai-race/id=118431/>.

¹² Michael Porter, *Clusters and the New Economics of Competition*, Harvard Business Review (Nov.-Dec. 1998), <https://hbr.org/1998/11/clusters-and-the-new-economics-of-competition>.

¹³ William R. Kerr & Frederic Robert-Nicoud, *Tech Clusters*, Journal of Economic Perspectives at 63 (Summer 2020), <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.34.3.50>.

¹⁴ Specifically, Seattle, Boston, San Francisco, San Diego, and San Jose. Robert D. Atkinson, et al., *The Case for Growth Centers: How to Spread Tech Innovation Across America*, Brookings (Dec. 9, 2019), <https://www.brookings.edu/research/growth-centers-how-to-spread-tech-innovation-across-america/>.

¹⁵ The NSF and other government agencies are doing admirable work, with the resources available, to encourage diverse research and create economies of scale for AI innovation, but they will not produce a strategic effect at the current level of effort, which is set against the backdrop of an overall decline in federal investment in R&D. Other notable recent federal initiatives include DARPA’s Artificial Intelligence Exploration Program, which fast tracks funding for awards up to \$1 million to explore feasibility of new AI concepts within an 18-month timeframe; and NSF’s National AI Research Institutes effort, which in 2020 funded seven multi-institution, university-based research institutes at \$4 million per year for five years and plans to launch another eight in 2021. *Artificial Intelligence at NSF*, NSF (Aug. 26, 2020), <https://www.nsf.gov/cise/ai.jsp>.

¹⁶ Funding for the National Institutes of Health (NIH) has grown from \$30 billion in 2010 to \$41 billion in 2020. *NIH Budget History: NIH Budget Mechanism Detail*, NIH Data Book (Oct. 2019), <https://report.nih.gov/nihdatabook/category/1; Budget>, NIH (June 29, 2020), <https://www.nih.gov/about-nih/what-we-do/budget>.

Chapter 11 - Endnotes

¹⁷ In 1953, the U.S. spent 0.72% of its GDP on R&D. In 1957, when the then-Soviet Union launched Sputnik, it had grown to 1.3%. R&D spending peaked at 1.86% in 1964. In 2017, it declined below 1953 levels to 0.61%. *Federal R&D Budget Dashboard*, American Association for the Advancement of Science (last accessed Jan. 14, 2021), <https://www.aaas.org/programs/r-d-budget-and-policy/federal-rd-budget-dashboard>.

¹⁸ The National AI Initiative Act of 2020 included in the National Defense Authorization Act for Fiscal Year 2021 creates a structure for a more strategic approach to harnessing AI through establishment of a National AI Initiative Office within the Office of Science and Technology Policy and associated advisory group and interagency construct. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁹ The NSF awarded grants for the first national AI research institutes in 2020, supporting seven university-based, multi-institution consortia organized around fundamental and applied areas of AI research, and plans to fund a second round of institutes in 2021, coordinating support not only with interagency partners but also with private sector stakeholders to launch eight additional institutes. *Artificial Intelligence at NSF*, NSF (Aug. 26, 2020), <https://www.nsf.gov/cise/ai.jsp>.

²⁰ Acting on a recommendation NSCAI issued in our *First Quarter Recommendations*, Congress has taken the first step to establish the NAIRR in the Fiscal Year 2021 National Defense Authorization Act, creating a task force to develop a roadmap for a future NAIRR. The result of this effort will be due to Congress 18 months after appointment of task force members. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021); see also *First Quarter Recommendations*, NSCAI at 12 (March 2020), <https://www.nscai.gov/previous-reports/>.


²¹ Since the explosion of deep learning in 2012 and accompanying growth in use of specialized hardware for AI computing, there has arisen what some have termed the “compute divide”—a disparity in access between large technology companies and elite universities and middle- and lower-tier universities to the resources necessary for cutting-edge AI research. Nur Ahmed & Muntasir Wahed, *The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, ArXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>.

²² This approach could build on successful models such as the COVID-19 High Performance Computing Consortium, (<https://covid19-hpc-consortium.org/>) and NSF's CloudBank, (<https://www.cloudbank.org/>).

²³ This would build on prior work undertaken by the Networking and Information Technology Research and Development (NITRD) Program Big Data Interagency Working Group. *Open Knowledge Network: Summary of the Big Data IWG Workshop*, National Science & Technology Council (Nov. 2018), <https://www.nitrd.gov/pubs/Open-Knowledge-Network-Workshop-Report-2018.pdf>. It would also build upon ongoing efforts through NSF's Convergence Accelerator track on Open Knowledge Networks. *NSF Convergence Accelerator Awards Bring Together Scientists, Businesses, Nonprofits to Benefit Workers*, NSF (Sept. 20, 2019), https://www.nsf.gov/news/special_reports/announcements/091019.jsp.

²⁴ Congress took an important step in the Consolidated Appropriations Act, 2021 by calling on the General Services Administration to create a five-year program to be known as the "AI Center of Excellence" (AI CoE) to "facilitate the adoption of artificial intelligence technologies in the Federal Government," among other duties. The AI CoE can help bridge discrete efforts across federal agencies to create a sizable market for government-specific AI applications. See *Rules Committee Print 116-68, Text of the House Amendment to Senate Amendment to H.R. 133*, U.S. House Committee on Rules at 378-81 (Dec. 11, 2020), <https://rules.house.gov/sites/democrats.rules.house.gov/files/BILLS-116HR133SA-RCP-116-68.pdf> (referring specifically to section 103 of the Consolidated Appropriations Act, 2021). In addition, the Defense Innovation Unit (DIU) is playing a role in creating markets at the intersection of AI and other strategic technologies through its project-based approach. Focus areas include AI applications for space systems, advanced diagnostics, semiconductors/advanced hardware, and other critical technologies identified by NSCAI in Chapter 16 of this report. DIU's experience indicates that creating a market for strategic technologies begins with the Department of Defense (DoD) and other government agencies pursuing an approach that is (a) contractually flexible, (b) aligned with firms' technological development plans, and (c) generating financial incentives through opportunities to scale production. *DIU Making Transformative Impact Five Years In*, DoD (Aug. 27, 2020), <https://www.defense.gov/Explore/News/Article/Article/2327021/diu-making-transformative-impact-five-years-in/>.

Chapter 12: Intellectual Property



Intellectual Property Policy is a National Security Priority



Issue Executive Order on IP for AI and Emerging Technologies.



Develop Plan to Reform and Establish IP Policies and Regimes.



Assess “IP Considerations.”



Propose Executive and Legislative Actions.




Integrate into National Security, Economic, and Technology Competitiveness Strategies.

China is both leveraging and exploiting intellectual property (IP) policies as a critical tool within its national strategies for emerging technologies. The United States has failed to similarly recognize the importance of IP in securing its own national security, economic interests, and technology competitiveness. The U.S. has not developed comprehensive IP policies to incentivize investments¹ in and protect the creation of artificial intelligence (AI) and other emerging technologies.² The consequence of this policy void—which includes legal uncertainties created by current U.S. patent eligibility and patentability doctrine, the lack of an effective response to China’s domestic and geopolitical strategies centered on its IP institutions,³ and the lack of effective data protection policies—is that the U.S. could lose its prime position in IP global leadership. At the same time, by strengthening its IP regimes,⁴ China is poised to “fill the void” left by weakened U.S. IP protections, particularly for patents, as the U.S. has lost its “comparative advantage in securing stable and effective property rights in new technological innovation.”⁵ This stark policy asymmetry has multiple significant domestic and international implications for the U.S.

First, U.S. courts have severely restricted what types of computer-implemented and biotech-related inventions can be protected under U.S. patent law.⁶ Critical AI and biotech-related inventions have been denied patent protection since 2010.⁷ Facing uncertainty in obtaining and retaining patent protection, inventors pursue trade secret protection. Trade secrets do not readily promote innovation markets, because trade secrets, unlike patents, do not contribute to accessible technical knowledge in the public domain.⁸ While these impacts might not be immediate, the long-term effects on AI and other emerging technology developments and competitiveness are concerning.⁹

Second, China has met its strategic policy goal of increasing the quantity of its patent applications and issued patents, creating the narrative that it has “won” the innovation race. In 2019, the total number of “invention” patent applications filed at the China National Intellectual Property Administration (CNIPA) was approximately three times as many as utility patent applications filed at the U.S. Patent and Trademark Office (USPTO).¹⁰ China also led the world in international patent applications under the Patent Cooperation Treaty (PCT) system of the World Intellectual Property Organization (WIPO).¹¹ Critically, China is now frequently identified as the current leader in domestic patent application filings for AI inventions.¹² Globally, AI patent applications originating from China outnumber those originating from the United States, especially in recent years.¹³



“The U.S. has not developed comprehensive IP policies to incentivize investments in and protect the creation of AI and other emerging technologies. The consequence of this policy void ... is that the U.S. could lose its prime position in IP global leadership.”

China's National IP Regimes for AI and Emerging Technologies

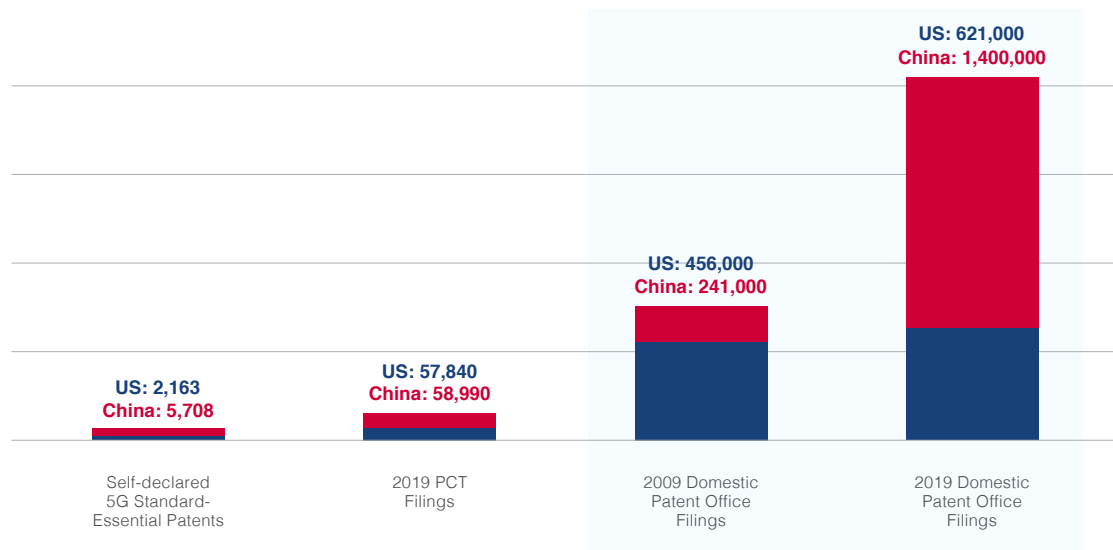
National 13th Five-Year Plan for the Development of Strategic Emerging Industries articulates IP-related goals for emerging technologies:

- Revising the Patent Law and Copyright Law
- Strengthening IP rights protections through rapid rights protection centers
- Developing strategic advancement plans for IP rights of emerging technologies
- Improving overseas IP rights and supporting companies involved in overseas M&A

Patent filings are incentivized by:

- Patent subsidies
- Rewards for granted patents
- Patent quotas set by provincial or municipal governments
- Preferential treatment in government procurement processes for companies with Chinese IP


Patent protection is increased through preliminary injunctions for patent infringement, increases in punitive damages for IP infringement (allows for quintuple damages for willful infringement), and specialized IP courts with efficient resolution and low litigation costs.



CSET Translation of *National 13th Five-Year Plan for the Development of Strategic Emerging Industries*, Central Committee of the Communist Party of China and the PRC State Council (Published Nov. 29, 2016) (translation by CSET on Dec. 9, 2019), <https://cset.georgetown.edu/research/national-13th-five-year-plan-for-the-development-of-strategic-emerging-industries/>; Eric Warner, *Patenting and Innovation in China: Incentives, Policy, and Outcomes*, RAND at 17-18 (Nov. 2014), <https://apps.dtic.mil/dtic/tr/fulltext/u2/a619128.pdf>; *Trademarks and Patents in China: The Impact of Non-Market Factors on Filing Trends and IP Systems*, U.S. Patent and Trade Office (Jan. 2021), <https://www.uspto.gov/sites/default/files/documents/USPTO-TrademarkPatentsInChina.pdf>; Ryan Davis, *4 Things to Know About China's Revised Patent Law*, Law 360 (Nov. 5, 2020), <https://www.law360.com/articles/1326419/>; Justice Tao Kaiyuan, *China's Commitment to Strengthening IP Judicial Protection and Creating a Bright Future for IP Rights*, WIPO Magazine (June 2019), https://www.wipo.int/wipo_magazine/en/2019/03/article_0004.html.

Note: The self-declared 5G standard-essential patent numbers are as of February 2020 and represent the combined total from the two companies that are the largest filers in each country. For the United States, 2,163 represents the 1,293 applications filed from Qualcomm and 870 from Intel. For China, 5,708 represents the 3,147 filed from Huawei and 2,561 filed from ZTE. This number also represents the standard-essential patents filed, not the number of patents granted. See Jed John Ikoba, *Huawei Has Filed the Most 5G Patents Globally as of February 2020 - Report*, Gizmochina (June 2, 2020), <https://www.gizmochina.com/2020/06/02/huawei-has-the-most-5g-standard-essential-patents-globally/>; *China Becomes Top Filer of International Patents in 2019 Amid Robust Growth for WIPO's IP Services, Treaties and Finances*, WIPO Media Center (Apr. 7, 2020), https://www.wipo.int/pressroom/en/articles/2020/article_0005.html; For the domestic patent office filings, according to the China National Intellectual Property Administration (CNIPA), "the number of invention patent applications it received increased by more than 500 percent between 2009 and 2019, from 241,000 to 1.4 million (although, interestingly, there was a 9 percent decrease from 2018 to 2019). In comparison the number of patent applications at the

USPTO increased by only 35% (from 456,000 to 621,000) over the same period. Hence, while in 2009 U.S. patent applications outnumbered Chinese applications by almost two-to-one, by 2019, the ratio had completely reversed. Most of the Chinese patenting increase can be attributed to applications filed by domestic applicants. Out of the 1.4 million CNIPA applications in 2019, domestic sources filed almost 90 percent (compared to 48 percent of USPTO applications).” See Patrick Thomas & Dewey Murdick, *Patents and Artificial Intelligence: A Primer*, Center for Security and Emerging Technology at 10 (Sept. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Patents-and-Artificial-Intelligence.pdf>. In 2019, there were also almost two million utility model applications in China. *Id.* at n. 17.



“China has met its strategic policy goal of increasing the quantity of its patent applications and issued patents, creating the narrative that it has “won” the innovation race.”


Third, regardless of quality concerns,¹⁴ China’s prolific patent application filings may further hurt U.S. innovators by creating a vast reservoir of “prior art” (the term in patent law for the worldwide scientific and technical knowledge by which an invention is evaluated to determine if it is new). This dramatically increases the quantity of prior art that must be reviewed in examining a patent application. As a result, the patent examination process at the USPTO will become increasingly difficult, if not onerous. At the same time, U.S. inventors may find it more difficult to obtain patents because they must show that their inventions are not disclosed in the prior art publications anywhere in the world, including in the Chinese-language patent applications filed in China and internationally.¹⁵ As Chinese patents come to dominate prior art searches by patent offices throughout the world, the current dominance of U.S. patents in worldwide prior art searches will erode.¹⁶

Fourth, and consistent with China’s extensive patent application filings, China’s companies have been identifying too many patents as “standard-essential” in standards development organizations, alleging that these patents must be practiced to comply with a technical standard.¹⁷ Although standard development organizations require patent owners to self-identify patents that may be deemed essential in future standards, these organizations leave final essentiality determinations to private companies negotiating licenses or, if there is a dispute, to courts.¹⁸ This practice of “overdeclaring” standard-essential patents (SEPs) furthers China’s global narrative that it has “won” the race to such standardized

technologies as 5G, prompting other countries to adopt China's technologies in their own communications infrastructures.¹⁹ A worrisome result may be that U.S. companies must pay billions in royalties to China's companies or face claims and resulting litigation that they willfully infringed on Chinese company patent rights.²⁰

Fifth, the lack of explicit legal protections for data or express policies on data ownership may hinder innovation and collaboration, particularly as technologies evolve.²¹ The absence of data protection regimes may disincentivize parties from making necessary investments to develop data sets that are critical for machine learning (ML) and AI systems.²² Additionally, the absence of data governance policies (such as contracting best practices) for IP-type protections or ownership rules could undermine the willingness of companies to enter into the public-private partnerships that are crucial for creating cutting-edge technological innovations.²³ This could also create challenges for U.S. collaboration with allies and other partners in vital AI R&D where data rights or ownership claims come into question.²⁴

Lastly, as further evidence that China views IP as essential in its domestic economic development, China continues to pervasively steal American IP-protected technological advances through varied means like cyber hacking of businesses and research institutes, technological espionage, blackmail, and illicit technology transfer.²⁵




“China continues to pervasively steal American IP-protected technological advances through varied means like cyber hacking of businesses and research institutes, technological espionage, blackmail, and illicit technology transfer.”

The IP Policy Void.

The U.S. Government needs to address these vulnerabilities resulting from the lack of comprehensive IP policies. Currently, the U.S. Government does not efficiently utilize IP policy as a tool to support national strategies for national security, economic interests, and technology competitiveness in AI and emerging technologies. The majority of the United States Government's coordinated IP policy efforts are focused on IP enforcement and preventing IP theft.²⁶ The U.S., however, lacks an agency or interagency entity that is empowered to both develop and execute national IP policies that support and integrate with national strategies. As a result, the United States lacks cohesive, legislatively mandated AI and emerging-technology IP policies that are integrated into national strategy frameworks to address, for example, global competition from countries like China.

America's IP laws and institutions must be considered as critical components for safeguarding U.S. national security interests, including advancing economic prosperity and technology competitiveness. The United States must, at a minimum, articulate and develop national IP reforms and policies with the goal of incentivizing, expanding, and protecting AI and emerging technologies, at home and abroad. Such policies should be developed and proposed via the Executive Branch with a process that integrates the disparate departments and agencies that serve important roles in promoting U.S. innovation. The Executive Branch should:



“America’s IP laws and institutions must be considered as critical components for safeguarding U.S. national security interests, including advancing economic prosperity and technology competitiveness.”

Recommendation

Develop and implement national IP policies to incentivize, expand, and protect AI and emerging technologies. The President should issue an executive order to recognize IP as a national priority and require the development of a comprehensive plan to reform and create IP policies and regimes that further national security, economic interests, and technology competitiveness strategies. The Commission recommends that the executive order direct the Vice President, as chair of the Technology Competitiveness Council (TCC), or otherwise as chair of an interagency task force, to oversee this effort. The executive order should direct the Secretary of Commerce—in coordination with the Under Secretary of Commerce for Intellectual Property and the Director of the USPTO²⁷—to develop proposals to reform and establish new IP policies and regimes, as needed, to incentivize, expand, and protect AI and emerging technologies. The plan should include proposals for executive and legislative actions for IP policy changes to achieve these objectives and should be accompanied by an assessment of a non-exhaustive list of “IP considerations.”²⁸ The Executive Order should direct the Vice President to assess which IP policies, regimes, and reform proposals from the Secretary of Commerce should be integrated into national security, economic, and technology competitiveness strategies and empower the Secretary of Commerce to facilitate implementation of such proposals.

National Intellectual
Property
Considerations.**Patent Eligibility****Combat IP Theft****Counter China's Narrative on "Winning" Tech Competition Based on Filings****Inventorship by AI****Impact of China's Filings on USPTO & U.S. Inventors****Global IP Alignment Efforts****Impediments to AI Public-Private Partnerships & International Collaboration****Democratize Innovation & IP Ecosystems****IP Protection for Data****"Standard-Essential" Patents Process**

Chapter 12 - Endnotes

¹ Advances in emerging technologies require significant investments. These investments are partly public, but advances also require extensive private investments.

² Technologies critical to national security interests include AI and biotechnology. NSCAI proposes an initial list of emerging technologies key to U.S. national competitiveness in Chapter 16 of this report.

³ CSET translation of *National 13th Five-Year Plan for the Development of Strategic Emerging Industries*, Central People's Government of the People's Republic of China at 59 (Nov. 29, 2016) (translation by CSET on Dec. 9, 2019), <https://cset.georgetown.edu/research/national-13th-five-year-plan-for-the-development-of-strategic-emerging-industries/>. China continues to make extensive reforms to its IP regimes in furtherance of its innovation and industrial competitiveness goals. See Mark Cohen, *IPO's Comments on Recent Patent Legislation: Untangling a Complex Web*, China IPR blog (Dec. 15, 2020), <https://chinaipr.com/2020/12/15/ipos-comments-on-recent-patent-legislation-untangling-a-complex-web/>.

⁴ China's actions include ensuring that AI and associated technologies are eligible for patent protection, increasing damages awards for patent infringement, continuing to issue preliminary injunctions for infringement of valid patents, and creating specialized IP courts with more efficient resolution of IP cases. See Kevin Madigan & Adam Mossoff, *Turning Gold into Lead: How Patent Eligibility Doctrine Is Undermining U.S. Leadership in Innovation*, *George Mason Law Review* at 943-946 (April 13, 2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2943431 [hereinafter *Turning Gold Into Lead*]; Ryan Davis, *4 Things to Know About China's Revised Patent Law*, *Law 360* (Nov. 5, 2020), <https://www.law360.com/articles/1326419/>; Liaoteng Wang et. al., *A Comparative Look at Patent Subject Matter Eligibility Standards: China Versus the United States*, *IP Watchdog* (June 12, 2020), <https://www.ipwatchdog.com/2020/06/12/comparative-look-patent-subject-matter-eligibility-standards-china-versus-united-states/id=122339/>; Erick Robinson, *Everything You Need to Know about China's New Preliminary Injunction Rules*, *IAM* (Dec. 21, 2018), <https://www.iam-media.com/designs/everything-you-need-know-about-chinas-new-preliminary-injunction-rules>; Justice Tao Kaiyuan, *China's Commitment to Strengthening IP Judicial Protection and Creating a Bright Future for IP Rights*, *World Intellectual Property Organization*, *WIPO Magazine* (June 2019), https://www.wipo.int/wipo_magazine/en/2019/03/article_0004.html.

⁵ See *Turning Gold Into Lead*, at 955.

⁶ See *Turning Gold Into Lead*. In January 2019, the United States Patent & Trademark Office (USPTO) published initial patent eligibility guidance that applies during examination of patent applications at the USPTO, which arguably decreased uncertainty as to patent eligibility determinations during the patent application examination and granting process. However, the United States Court of Appeals for the Federal Circuit, the appellate court with jurisdiction of appeals from patent cases, held that it is not bound by the Guidance. See *Cleveland Clinic Found. v. True Health Diagnostics LLC*, 760 F. App'x at 1013, 1020 (Fed. Cir. 2019) (non-precedential); *In re Rudy*, 956 F.3d 1379, 1383 (Fed. Cir. 2020) (precedential) (citing *Cleveland Clinic Found.*, 760 F. App'x at 1021).

⁷ *Athena Diagnostics v. Mayo Collaborative Services*, 915 F.3d 743 (Fed. Cir. 2019), rehearing en banc denied 927 F.3d 1333 (Fed. Cir. 2019) (method of diagnosing certain, previously undiagnosable, patients suffering from the neurological disorder myasthenia gravis using MuSK autoantibodies); *The Cleveland Clinic Found. v. True Health Diagnostics LLC*, 760 F. App'x 1013 (Fed. Cir. 2019) (method of assessing the risk a patient has cardiovascular disease by analyzing the level of a certain enzyme in a patient's blood); *Roche Molecular Systems, Inc. v. Cepheid*, 905 F.3d 1363 (Fed. Cir. 2018) (DNA primers used in a method to detect the pathogenic bacterium *Mycobacterium tuberculosis*); *Ariosa Diagnostics, Inc. v. Sequenom, Inc.*, 788 F.3d 1371 (Fed. Cir. 2015), *cert. denied*, 136 S. Ct. 2511 (2016) (method of diagnosing fetal characteristics based on paternally inherited DNA found in a mother's bloodstream without creating a major health risk for the fetus); *PUREPREDICTIVE, Inc. v. H2O.AI, Inc.*, No. 17-cv-03049-WHO, 2017 WL 3721480 (N.D. Cal. Aug. 29, 2017) (predictive analytics); *Power Analytics Corp. v. Operation Tech., Inc.*, No. 16-cv-01955 JAK (FFMx), 2017 WL 5468179 (C.D. Cal. July 13, 2017) ("computer simulation techniques with real-time system monitoring and prediction of electrical system performance").

Chapter 12 - Endnotes

⁸ See *Crash Course on Patents: What Is a Patent and Why Is It Useful*, Ius Mentis (last accessed Dec. 30, 2020), <https://www.iusmentis.com/patents/crashcourse/whatis/> (because patents openly publish details of the invention, other inventors can license this invention or think of enhancements or design around the disclosure); Steven Hoffman & Calla Simeone, *Trade Secret Protection & the COVID-19 Cure: Observations on Federal Policy-Making & Potential Impact on Biomedical Advances*, JDSupra (Sept. 15, 2020), <https://www.jdsupra.com/legalnews/trade-secret-protection-the-covid-19-37383/> (discussing implications of uncertainty in patent eligibility on use of trade secrets for biomedical advances).

⁹ Surveys and industry reports demonstrate that investment has already shifted away from patent-intensive industries. See Mark F. Schultz, *The Importance of an Effective and Reliable Patent System to Investment in Critical Technologies*, Alliance for U.S. Startups and Investors for Jobs at 24-37 (July 2020), https://static1.squarespace.com/static/5746149f86db43995675b6bb/t/5f2829980ddf0c536e7132a4/1596467617939/USIJ+Full+Report_Final_2020.pdf.

¹⁰ Patrick Thomas & Dewey Murdick, *Patents and Artificial Intelligence: A Primer*, Center for Security and Emerging Technology at 10 (Sept. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Patents-and-Artificial-Intelligence.pdf> [hereinafter CSET, A Primer]; *U.S. Patent Statistics Chart Calendar Years 1963-2019*, USPTO (April 2020), https://www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.htm.

¹¹ See CSET, A Primer at 11; Aaron Winger, *China Surpasses U.S. to Become Top Filer of PCT International Patent Applications in 2019*, National Law Review (April 7, 2020), <https://www.natlawreview.com/article/china-surpasses-us-to-become-top-filer-pct-international-patent-applications-2019>. China is on pace to continue being the top PCT filer in 2020. See Aaron Winger, *China 2020 H1 Patent Data Indicates China Likely to Remain Top International Filer in 2020*, National Law Review (July 11, 2020), <https://www.natlawreview.com/article/china-2020-h1-patent-data-indicates-china-likely-to-remain-top-international-filer>.

¹² *AI Innovators*, RS (last accessed Dec. 30, 2020), <https://uk.rs-online.com/web/generalDisplay.html?id=did-you-know/ai-innovators>; George Leopold, *China Dominates AI Patent Filings*, Enterprise AI (Aug. 31, 2020), <https://www.enterpriseai.news/2020/08/31/china-dominates-ai-patent-filings/>; CSET, A Primer.

¹³ CSET, A Primer at 9, 12, n. 23.

¹⁴ *Trademarks and Patents in China: The Impact of Non-Market Factors on Filing Trends and IP Systems*, USPTO at 1 (Jan. 2021), <https://www.uspto.gov/sites/default/files/documents/USPTO-TrademarkPatentsInChina.pdf>; Jonathan Putnam, et al., *Innovative Output in China*, SSRN at 32 (Aug. 2020) (pending revision), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3760816.

¹⁵ Jeanne Suchodolski, et al., *Innovation Warfare*, North Carolina Journal of Law & Technology at 201 (Dec. 7, 2020), <https://ncjolt.org/articles/volume-22/volume-22-issue-2/innovation-warfare/> [hereinafter *Innovation Warfare*].

¹⁶ Rob Sterne, *How China Will Fundamentally Change the Global IP System*, IP Watchdog (July 24, 2019), <https://www.ipwatchdog.com/2019/07/24/china-changing-global-ip-system/id=111613/>.

¹⁷ Over-declaration is already present in 5G. See Matthew Noble, et al., *Determining Which Companies Are Leading the 5G Race*, IAM (July/August 2019), <https://www.twobirds.com/~media/pdfs/news/articles/2019/determining-which-companies-are-leading-the-5g-race.pdf?la=en&hash=8ABA5A7173EEE8FFA612E070C0EA4B4F53CC50DE>; *Meeting the China Challenge: A New American Strategy for Technology Competition*, Working Group on Science and Technology in U.S.-China Relations at 27, 29 (Nov. 16, 2020), https://china.ucsd.edu/files/meeting-the-china-challenge_2020_report.pdf [hereinafter *Meeting the China Challenge*].

¹⁸ *IEEE SA Standards Board Bylaws*, IEEE Standards Association (last accessed Jan. 15, 2020), <https://standards.ieee.org/about/policies/bylaws/sect6-7.html#loa>.

¹⁹ *5G Technological Leadership*, Hudson Institute (Dec. 5, 2020), <https://www.hudson.org/research/16547-5-g-technological-leadership>; *Innovation Warfare*, at 201, n.130 (China's firms recognize the strategic importance of standard-setting activities and that participation in those forums provides the legal means to both access and influence developing technologies).

²⁰ Because standard-essential patents (SEPs) may reach into the hundreds of thousands for technologies, licensing fees carry significant economic repercussions. See *5G Technological Leadership*, Hudson Institute at 3 (Dec. 5, 2020), <https://www.hudson.org/research/16547-5-g-technological-leadership> (“[P]atent counting might have negative consequences on firms working in the US innovation economy ... if judges or regulators rely on simple counts of total patents as a metric for determining the value of patent portfolios. The failure to account for differences in patent quality risks overcompensating some patent holders, namely those with less valuable technologies, but undercompensating those that have developed breakthrough innovation.”); Andrei Iancu, Director of USPTO, Remarks at the Center for The Protection of Intellectual Property 2020 Fall Conference (Oct. 7, 2020), <https://cpip.gmu.edu/2020/10/20/cpip-2020-fall-conference-day-one-recap/>; Muzammil Hassan, et al., *Who Owns Core 5G Patents? A Detailed Analysis of 5G SEPs*, GreyB (2020), <https://www.greyb.com/5g-patents/#The-State-of-Declared-5G-Patents>; Cody M. Akins, *Overdeclaration of Standard-Essential Patents*, Texas Law Review (2020), https://texaslawreview.org/wp-content/uploads/2020/02/Akins_Printer.pdf.

²¹ Mitchell Smith, *A Comparison of the Legal Protection of Databases in the United States and EU: Implications for Scientific Research*, SSRN (May 23, 2010), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1613451; Daniel J. Gervais, *Exploring the Interfaces Between Big Data and Intellectual Property Law*, *Journal of Intellectual Property, Information Technology and Electronic Commerce Law* (2019), <https://scholarship.law.vanderbilt.edu/faculty-publications/1095>.

²² In the USPTO report surveying stakeholders for perspectives on IP policy for AI, “[c]ommenters were nearly equally divided between the view that new intellectual property rights were necessary to address AI inventions and the belief that the current U.S. IP framework was adequate to address AI inventions. Generally, however, commenters who did not see the need for new forms of IP rights suggested that developments in AI technology should be monitored to ensure needs were keeping pace with AI technology developments. The majority of opinions requesting new IP rights focused on the need to protect the data associated with AI, particularly ML.” *Public Views on Artificial Intelligence and Intellectual Property Policy*, USPTO at 15 (Oct. 2020), https://www.uspto.gov/sites/default/files/documents/USPTO_AI-Report_2020-10-07.pdf.

²³ Thomas E. Ayers, *Changing How We Buy Weapons Will Benefit Industry, Government and Taxpayers*, *Defense News* (Nov. 20, 2019), <https://www.defensenews.com/opinion/commentary/2019/11/20/changing-how-we-buy-weapons-will-benefit-industry-government-and-taxpayers/> (discussing the tension between Air Force and vendors over IP protection).

²⁴ See also the Chapter 15 Blueprint for Action.

²⁵ Meeting the China Challenge, at 4, 16.

²⁶ *Annual Intellectual Property Report to Congress*, U.S. Intellectual Property Enforcement Coordinator (March 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/04/IPEC-2019-Annual-Intellectual-Property-Report.pdf> (providing an overview of IP responsibilities across the United States government).

²⁷ Other Executive Branch departments and agencies, and the U.S. Copyright Office, should resource and support the Secretary of Commerce in these efforts.

²⁸ A non-exhaustive list of IP considerations should include patent eligibility doctrine, countering China's narrative on “winning” AI innovation based on patent application filings, the impact of China's patent application filings on USPTO's examination process and U.S. inventors, impediments in IP contractual system to public-private partnerships and international collaboration, IP protections for data, combatting IP theft, AI inventorship, global IP alignment, democratizing innovation and IP ecosystems, and SEPs process.

Chapter 13: Microelectronics



Rebuilding U.S. Microelectronics Leadership



Stay Two Generations Ahead in Microelectronics



Multiple Sources of Domestic Cutting-Edge Manufacturing



National Microelectronics Strategy




Double Down on Microelectronics R&D



Tax Credits and Grants for U.S. Fabrication Facilities

U.S. leadership in microelectronics is critical to overall U.S. leadership in artificial intelligence (AI). Several assessments underpin this argument:

- Hardware is a foundational element of the AI stack alongside data, algorithms, and talent.¹
- Exponential increases in computational power have driven the last decade of progress in machine learning (ML).²
- After decades leading the microelectronics industry, the United States will soon source roughly 90% of all high-volume, leading-edge integrated-circuit production from countries in East Asia.³ This means the United States is almost entirely reliant on foreign sources for production of the cutting-edge semiconductors critical for defense systems and industry more broadly, leaving the U.S. supply chain vulnerable to disruption by foreign government action or natural disaster.
- Specialized hardware, novel packaging techniques such as heterogeneous integration and 3D stacking, and new types of devices will drive future AI developments as traditional architectures of silicon-based chipsets encounter diminishing marginal performance improvements.⁴
- Demand for trusted microelectronics will only grow as the military and Intelligence Community (IC) continue to incorporate AI into mission-critical systems.⁵



“... the United States is almost entirely reliant on foreign sources for production of the cutting-edge semiconductors critical for defense systems and industry more broadly, leaving the U.S. supply chain vulnerable to disruption by foreign government action or natural disaster.”

U.S. leadership in semiconductors has long been taken for granted based on America's advantage as a pioneer of the microelectronics industry. Gradually, however, the United States has been losing its edge. Although American universities and firms remain global leaders in the key areas of semiconductor R&D and chip design, the semiconductor industry is now highly globalized and competitive. Taiwan Semiconductor Manufacturing Corporation (TSMC) leads the world in semiconductor contract manufacturing, and Samsung in South Korea is also producing state-of-the-art logic chips.⁶ TSMC also leads in the production of ARM-based chips, which is becoming the predominant chip architecture for mobile devices, servers, and other key applications of emerging technologies.⁷ In a bid to catch up and achieve chip self-sufficiency, China is pursuing unprecedented state-funded efforts to forge a world-leading semiconductor industry by 2030. Although China is behind firms headquartered in Taiwan, South Korea, and the U.S. in terms of chip manufacturing, it is advancing quickly.⁸ Meanwhile, Intel, the leading U.S. manufacturer, remains competitive in chip design but has faced manufacturing setbacks for leading-edge chips and may fall further behind its rivals in Taiwan and South Korea. Current projections put the firm two generations or more behind the cutting-edge node by 2022.⁹ These and other concerning trends indicate that America's leadership in microelectronics is eroding, especially in manufacturing, assembly, testing, and packaging.¹⁰


The dependency of the United States on semiconductor imports, particularly from Taiwan, creates a strategic vulnerability for both its economy and military to adverse foreign government action, natural disaster, and other events that can disrupt the supply chains for electronics. Despite tremendous expertise in microelectronics research, development, and innovation across the country, the United States is constrained by a lack of domestically-located semiconductor fabrication facilities, especially for state-of-the-art semiconductors. If current trends continue, the United States will soon be unable to catch up in fabrication, and could eventually also be outpaced in microelectronics design. If a potential adversary bests the United States in semiconductors over the long term or suddenly cuts off U.S. access to cutting-edge chips entirely, it could gain the upper hand in every domain of warfare. Focusing the efforts of the U.S. Government, industry, and academia to develop domestic microelectronics fabrication facilities will reduce dependence on imports, preserve leadership in technological innovation, support job creation, improve national security and balance of trade, and enhance the technological superiority and readiness of the military, which is an important consumer of advanced microelectronics.



“Despite tremendous expertise in microelectronics research, development, and innovation across the country, the United States is limited by a lack of domestically-located semiconductor fabrication facilities...”

State-of-the-Art Semiconductor Manufacturing by Firm: 2014-2024.

State-of-the-Art Semiconductor Manufacturing by Firm: 2014-2024



Node size for 2021-2024 are projections and reflect firm roadmaps

Node size reflects estimated first year of mass production

No roadmap displayed beyond 2021 for SMIC due to export control restrictions on materials currently required for production beyond 14 nm

To regain U.S. leadership in microelectronics, the Executive Branch should finalize and implement a national microelectronics leadership strategy. Additionally, Congress should create a 40% refundable tax credit for domestic fabrication investments by firms from the United States and its allies and appropriate an additional \$12 billion over the next five years for microelectronics research, development, and infrastructure. Together these efforts will enable the U.S. government, private sector, and academia to rise to the challenge of rebuilding U.S. semiconductor superiority.

Objective: Stay two generations ahead of China in state-of-the-art microelectronics and maintain multiple sources of cutting-edge microelectronics fabrication in the United States.

The United States should focus the attention and resources necessary for long-term competition in microelectronics by adopting an overarching national objective: to stay two generations ahead of potential adversaries in state-of-the-art microelectronics while also maintaining multiple sources of cutting-edge microelectronics fabrication inside the United States.¹¹ While the United States has historically led China by at least two generations in semiconductor design and fabrication, this has not been an explicit policy goal. And while China has not been able to surpass the United States, other nations such as Taiwan and South Korea now clearly lead the U.S. in state-of-the-art semiconductor manufacturing. This leaves the U.S. reliant on foreign sources for critical inputs to defense systems and U.S. industry more broadly. Yet the United States retains a strong position in segments of the global value chain for semiconductors, such as design, electronic design automation tools, and semiconductor manufacturing equipment (SME).¹² Therefore, an objective to rebuild microelectronics leadership should be stated plainly to concentrate national support across government, industry, and academia on regaining leadership in sectors such as semiconductor fabrication where the United States has fallen behind and also to track progress over time against a clear yardstick. To achieve this objective, the Commission recommends focusing action along three fronts:

- Implementing a national microelectronics strategy;
- Revitalizing domestic microelectronics fabrication by incentivizing multiple cutting-edge domestic fabrication facilities; and
- Ramping up microelectronics research.

In addition to these efforts to promote U.S. microelectronics leadership, the United States and its allies should utilize targeted export controls on high-end semiconductor manufacturing equipment, described in Chapter 14 of this report, to protect existing technical advantages and slow the advancement of China's semiconductor industry.

U.S. leadership in microelectronics is essential to overall U.S. leadership in AI.




Recommendation

Implement the National Microelectronics Strategy. The United States lacks a national microelectronics strategy to coordinate semiconductor policy, funding, and incentives within the Executive Branch and externally with industry and academia. A truly national strategy would build on this Commission's work as well as previous studies conducted by the United States government or on its behalf. It would also integrate the disparate approaches of the Departments of State, Defense, Energy, Commerce, and Treasury, and other relevant agencies, to promote domestic R&D and semiconductor manufacturing expertise while preventing the illicit transfer of technology to competitors. Finally, it would be updated on a consistent basis to foster a coordinated approach and adapt to shifting challenges to microelectronics innovation, competitiveness, and supply chain integrity.

In line with the Commission's recommendations, the Fiscal Year 2021 National Defense Authorization Act (NDAA) creates a subcommittee of the National Science and Technology Council (NSTC), consisting of senior government officials, to develop a National Strategy on Microelectronics Research and oversee its implementation.¹³ However, for this key effort to be successful, it should be prioritized by the White House by requiring the NSTC subcommittee to submit the National Microelectronics Strategy to the President within 270 days.

Recommendation


Revitalize domestic microelectronics fabrication. The Commission concludes that the United States is overly dependent upon globally diversified supply chains for microelectronics, including imports from potential adversaries. Furthermore, as a result of gaps in the U.S. industrial base, the risks are increasing that the United States could lose access to trusted, assured, and state-of-the-art semiconductors for national security use cases. Despite these concerns, the Commission has been encouraged by a number of developments over the past year to revitalize the domestic fabrication of state-of-the-art microelectronics.



“... the United States could lose access to trusted, assured, and state-of-the-art semiconductors for national security use cases.”

Examples include TSMC’s decision to develop an advanced facility in the United States and Intel’s publicly stated interest in working with the United States government to develop a commercial U.S. foundry.¹⁴ However, these are only initial steps, and more must be done by the U.S. government to reach an end state where multiple firms are fabricating state-of-the-art chips domestically. Without several U.S.-based fabrication facilities, both U.S. industry and U.S. national security face risks from competitive pressures and supply chain shortages. The most significant recent development has been the inclusion of several semiconductor-related provisions from the “CHIPS for America Act” in the Fiscal Year 2021 National Defense Authorization Act (NDAA).¹⁵ However, these programs require sufficient appropriations to succeed, and they did not receive appropriated funding in Fiscal Year 2021, which leaves congressional priorities unclear. Further congressional action to establish refundable investment tax credits and set the conditions for the domestic production of advanced microelectronics will be important to enable the United States to remain two generations ahead of China. Specifically, the U.S. government should:

- **Incentivize domestic leading-edge merchant fabrication through refundable investment tax credits.** Although introduced as part of the CHIPS for America Act, Congress has not yet passed legislation establishing a 40% refundable investment tax credit for semiconductor facilities and equipment.¹⁶ Existing U.S. incentives reduce the cost of foundry construction attributable to capital expenses, operating expenses, and taxes by just 10% to 15%. A credit of this magnitude is needed to make the United States a competitive market for semiconductor manufacturing, as other leading semiconductor manufacturing nations such as South Korea, Taiwan, and Singapore offer 25% to 30% cost reduction, roughly double what the United States currently offers.¹⁷ This gap in incentives is one driving factor behind the lack of an advanced logic merchant foundry in the United States. Closing the incentive gap and broadening it to include companies from allied countries will incentivize U.S. firms to construct facilities domestically while also attracting foreign firms such as TSMC and Samsung. Additionally, increasing demand in the United States for high-end SME will create new business opportunities for SME manufacturers from allied countries, particularly Japan and the Netherlands, which could increase their governments’ willingness to align their export control policies with strict U.S. policies prohibiting the export of such equipment to China.¹⁸



“... other leading semiconductor manufacturing nations such as South Korea, Taiwan, and Singapore offer 25 to 30 percent cost reduction, roughly double what the United States currently offers.”

Recommendation

Double-down on federally funded microelectronics research. Each succeeding generation of chips using traditional architectures of silicon-based transistors faces diminishing marginal gains to performance as they reach the limits imposed by the laws of physics. As a result, the relative advantage the United States has enjoyed by staying roughly two generations ahead of potential adversaries in the design phase of developing cutting-edge hardware could decrease over time as the gap between hardware generations narrows. Therefore, the United States must look to heterogeneous integration and other novel hardware improvements in the medium term to continue out-innovating competitors. Over the longer term, the United States must also continue its portfolio approach to future microelectronics pathways by investing in new materials and entirely new hardware approaches, such as quantum and neuromorphic computing. Broad-based investments and incentives will also be important to maintain leadership in other areas of U.S. strength related to semiconductor manufacturing, including electronic design automation tools and SME.

Four primary research arms of the United States government focused on both medium- and long-term microelectronics breakthroughs are the Department of Energy, Defense Advanced Research Projects Agency (DARPA), National Science Foundation (NSF), and the Department of Commerce, primarily through engagement with industry. Their suite of existing programs, such as DARPA's Electronics Resurgence Initiative, is targeting the right research areas but must expand by an order of magnitude to achieve the necessary breakthroughs and maintain U.S. competitiveness. Additional funding should support not only research projects, but also the capital-intensive underlying infrastructure for microelectronics development, including the National Semiconductor Technology Center and advanced packaging prototyping activities authorized in the Fiscal Year 2021 NDAA. In particular, advances in packaging will be critical to future improvements in semiconductor

capabilities as firms reach physical limits for two-dimensional transistor density.¹⁹ The government should:

- **Double down on federal research funding to lead the next generation of microelectronics.** The Commission recommends substantially increasing the United States government's full range of research efforts focused on microelectronics. Congress should appropriate an additional \$1.1 billion for semiconductor research and \$1 billion for the Advanced Packaging National Manufacturing Program in Fiscal Year 2022. Building on these investments, these funding levels should continue for five years, for a total investment of roughly \$12 billion. These amounts are consistent with the funding levels introduced, but not yet appropriated, in the CHIPS for America Act²⁰ and the American Foundries Act of 2020.²¹ In line with the existing focus areas of these programs and the Commission's prior recommendations, the funding should be applied to developing infrastructure and pursuing breakthroughs in promising areas such as next-generation tools beyond extreme ultraviolet lithography, 3D chip stacking, photonics, carbon nanotubes, gallium nitride transistors, domain-specific hardware architectures, electronic design automation, and cryogenic computing.



“... advances in packaging will be critical to future improvements in semiconductor capabilities as firms reach physical limits for two-dimensional transistor density.”

Chapter 13 - Endnotes

¹ Dave Martinez, et al., *Artificial Intelligence: Short History, Present Developments, and Future Outlook*, MIT at 27, n. 10 (Jan. 2019), <https://www.ll.mit.edu/sites/default/files/publication/doc/2019-09/Artificial%20Intelligence%20Short%20History%2C%20Present%20Developments%2C%20and%20Future%20Outlook%20-%20Final%20Report%20-%20Martinez.pdf> (citing Andrew Moore, et al., *The AI Stack: A Blueprint for Developing and Deploying Artificial Intelligence*, International Society for Optics and Photonics, Ground/Air Multisensor Interoperability, Integration, and Networking for Persistent ISR IX [2018]).

² Recent machine learning (ML) breakthroughs have relied heavily on computing power, and the amount of compute used in the largest AI training runs has been increasing exponentially since 2012. Girish Sastry, et al., *Addendum: Compute Used in Older Headline Results*, OpenAI (Nov. 7, 2019), <https://openai.com/blog/ai-and-compute/#addendum>.

³ Michaela Platzer, et al., *Semiconductors: U.S. Industry, Global Competition, and Federal Policy*, Congressional Research Service at 12 (Oct. 26, 2020), <https://crsreports.congress.gov/product/pdf/R/R46581>.

⁴ Sara Hooker, *The Hardware Lottery*, arXiv (Sept. 21, 2020), <https://arxiv.org/pdf/2009.06489.pdf>.

⁵ Gaurav Batra, et al., *Artificial Intelligence Hardware: New Opportunities for Semiconductor Companies*, McKinsey & Co. (Jan. 2, 2019), <https://www.mckinsey.com/industries/semiconductors/our-insights/artificial-intelligence-hardware-new-opportunities-for-semiconductor-companies>.

⁶ Taiwan Semiconductor Manufacturing Corporation (TSMC) has already begun producing 5nm state-of-the-art logic chips and aims to produce 3nm chips by the end of 2021. Samsung is also producing 5nm chips. Intel does not anticipate producing 7nm chips in-house until at least 2022 and may outsource manufacturing to TSMC. Firms in China are producing 12 nm chips. Richard Waters, *Intel Looks to New Chief's Technical Skills to Plot Rebound*, Financial Times (Jan. 14, 2021), <https://www.ft.com/content/51f63b07-aeb8-4961-9ce9-c1f7a4e326f0>; Mark Lapedus, *China Speeds Up Advanced Chip Development*, Semiconductor Engineering (June 22, 2020), <https://semiengineering.com/china-speeds-up-advanced-chip-development/>; *5nm Technology*, TSMC (last accessed Jan. 16, 2021), https://www.tsmc.com/english/dedicatedFoundry/technology/logic/l_5nm; Debby Wu, *TSMC's \$28 Billion Spending Blitz Ignites a Global Chip Rally*, Bloomberg (Jan. 14, 2021), <https://www.bloomberg.com/news/articles/2021-01-14/tsmc-profit-beats-expectations-as-chipmaker-widens-tech-lead>; Anton Shilov, *Samsung Foundry Update: 5nm SoCs in Production, HPC Shipments to Expand in Q4*, Tom's Hardware (Nov. 1, 2020), <https://www.tomshardware.com/news/samsung-foundry-update-5nm-socs-in-production-hpc-shipments-to-expand-in-q4>.

⁷ ARM and TSMC Announce Multi-Year Agreement to Collaborate on 7nm FinFET Process Technology for High-Performance Compute, Design & Reuse (March 15, 2016), <https://www.design-reuse.com/news/39433/arm-tsmc-7nm-finfet.html>.

⁸ Michaela D. Platzer, et al., *Semiconductors: U.S. Industry, Global Competition, and Federal Policy*, Congressional Research Service at 2, 25, 27 (Oct. 26, 2020), <https://crsreports.congress.gov/product/pdf/R/R46581>.

⁹ Ian King, *Intel 'Stunning Failure' Heralds End of Era for U.S. Chip Sector*, Bloomberg (July 24, 2020), <https://www.bloomberg.com/news/articles/2020-07-25/intel-stunning-failure-heralds-end-of-era-for-u-s-chip-sector>.

¹⁰ Michaela D. Platzer, et al., *Semiconductors: U.S. Industry, Global Competition, and Federal Policy*, Congressional Research Service (Oct. 26, 2020), <https://crsreports.congress.gov/product/pdf/R/R46581>.

¹¹ The Commission's previous reports offered a range of initial recommendations to expand access to trusted semiconductors, increase microelectronics R&D funding, control the export of high-end semiconductor manufacturing equipment to adversaries, and reshore leading-edge fabrication facilities.

¹² John VerWey, *The Health and Competitiveness of the U.S. Semiconductor Manufacturing Equipment Industry*, U.S. International Trade Commission: Office of Industries Working Paper (July 1, 2019), <http://dx.doi.org/10.2139/ssrn.3413951>.

¹³ See Pub. L. 116-283, sec. 9906, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁴ Stephen Nellis, *Phoenix Okays Development Deal with TSMC for \$12 Billion Chip Factory*, Reuters (Nov. 18, 2020), <https://www.reuters.com/article/us-tsmc-arizona/phoenix-okays-development-deal-with-tsmc-for-12-billion-chip-factory-idUSKBN27Y30E>; Asa Fitch, et al., *Trump and Chip Makers Including Intel Seek Semiconductor Self-Sufficiency*, Wall Street Journal (May 11, 2020), <https://www.wsj.com/articles/trump-and-chip-makers-including-intel-seek-semiconductor-self-sufficiency-11589103002>.

¹⁵ See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021). These provisions authorize several programs the Commission has previously identified as essential to U.S. microelectronics leadership. In particular, the provisions would require drafting a National Microelectronics Leadership Strategy, establishing a National Semiconductor Technology Center, and creating an incubator for semiconductor startup firms and an Advanced Packaging National Manufacturing Institute, all of which align with previous recommendations from the Commission.

¹⁶ This incentive would reduce a semiconductor firm's tax bill by 40% on semiconductor manufacturing equipment and facilities through 2024, followed by reduced tax credit rates of 30% and 20% respectively, through 2025 and 2026

¹⁷ Antonio Varas, et al., *Government Incentives and US Competitiveness in Semiconductor Manufacturing*, BCG and SIA (Sept. 2020), <https://web-assets.bcg.com/27/cf/9fa28eeb43649ef8674fe764726d/bcg-government-incentives-and-us-competitiveness-in-semiconductor-manufacturing-sep-2020.pdf>.

¹⁸ See Chapter 14 of this report for additional details on export controls on SME.

¹⁹ *Heterogeneous Integration Roadmap: Chapter 1: HIR Overview and Executive Summary*, IEEE Electronics Packaging Society (Oct. 2019), https://eps.ieee.org/images/files/HIR_2019/HIR1_ch01_overview.pdf.

²⁰ See S. 3933 and H.R. 7178, Creating Helpful Incentives to Produce Semiconductors (CHIPS) for America Act, 116th Congress (2020).

²¹ See S. 4130, American Foundries Act of 2020, 116th Congress (2020).

Chapter 14: Technology Protection

Improving U.S. Technology Protection Capabilities



Enhance
Regulatory
Capacity



Utilize Targeted
Export Controls



Increase
Investment
Screening
Disclosures



Strengthen
Research
Protections



Preserve U.S.
Innovation
Advantages

America's ability to out-innovate competitors is the dominant component of any U.S. strategy for technology leadership. Promoting research, entrepreneurship, and talent development remain the key ingredients of success. However, as the margin of U.S. technological advantage narrows and foreign efforts to acquire American know-how and technology increase, the United States must also reexamine how it can protect ideas, hardware, companies, and its values.

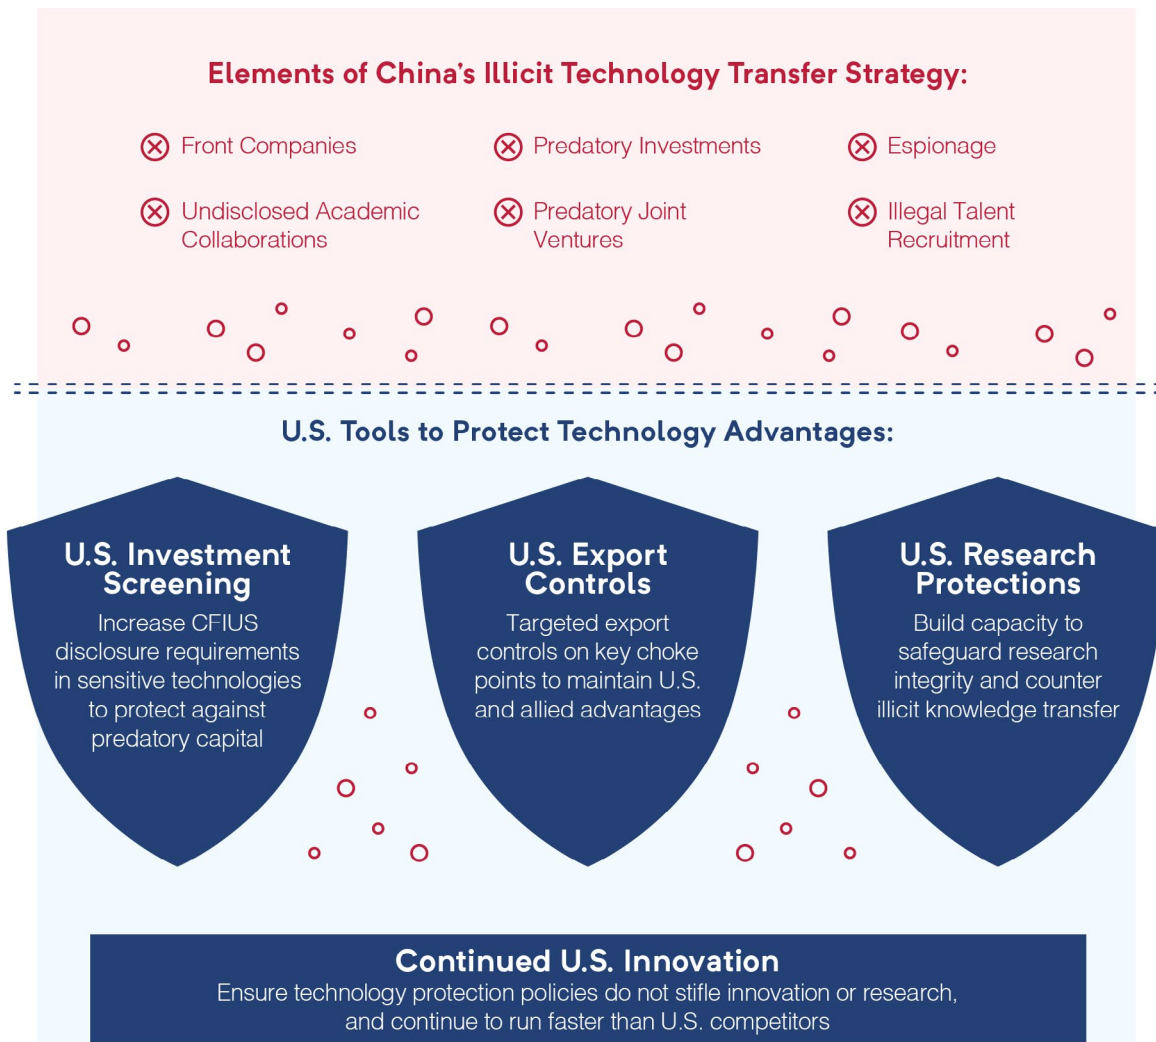
The United States confronts sustained threats from state-directed technology transfer and theft targeting artificial intelligence (AI) and other cutting-edge, dual-use technologies and basic research. China poses the most significant challenge. The Chinese Communist Party (CCP) has embarked on a multi-pronged campaign of licit and illicit technology transfer to become a "science and technology superpower" by 2050.¹ The campaign deliberately targets U.S. critical sectors, companies, and research institutions.² China's theft of U.S. technology—be it through circumventing export controls, commercial deals with U.S. companies to access intellectual property (IP), or espionage—costs the United States \$300 billion to \$600 billion per year.³ And that only captures immediate losses, not the ongoing damage to the U.S. economy over time. China simultaneously exploits open research environments through cyber-enabled intrusion, talent recruitment programs, and manipulated research partnerships.⁴ In effect, China is using American taxpayers' dollars to fund its military and economic modernization.



“China’s theft of U.S. technology—be it through circumventing export controls, commercial deals with U.S. companies to access intellectual property, or espionage—costs the United States \$300–\$600 billion per year.”

Russia also poses a significant illicit technology transfer threat, particularly as it relates to technologies with defense applications. Although the Russian government’s efforts to steal U.S. technology and IP do not operate at the same scale of comparable CCP efforts, Russia nonetheless is an aggressive and capable collector of technologies. It is likely to pose continued technology transfer threats over the coming decade, utilizing existing commercial and academic entities, as well as traditional and cyber espionage.⁵


Protect.



Modernizing Export Controls and Investment Screening.

How the United States designs policies to limit the movement of commercial goods or capital in the interests of national security will be one of the defining challenges of the next decade, as dual-use commercial technologies become increasingly important to national security. Export controls can and should be utilized not only to prevent the transfer of particularly sensitive equipment to strategic competitors, but also to slow competitors’ efforts to develop indigenous industries in sensitive technologies with defense applications.

If executed properly, export controls that slow competitors can sustain existing U.S. defense advantages over long periods of time. For instance, U.S. export controls on jet engine technology have stymied Chinese government-led efforts to produce a modern jet engine domestically for use in military aircraft for nearly 30 years.⁶



“How the United States designs policies to limit the movement of commercial goods or capital in the interests of national security will be one of the defining challenges of the next decade ...”

However, as currently designed and utilized, U.S. export controls and investment screening procedures are imperfect instruments for the AI competition. As a policy matter, investment screening and export controls were designed for a different era, when the distinction between civil and military technologies was clearer and there was little overlap between the economies of the United States and its competitors. Both conditions have changed. AI is dual-use and the emerging technology economies of the United States and China are deeply interconnected, which makes it extremely difficult to design controls that are feasible, maximize strategic impact, and minimize economic costs. While these tradeoffs are not new, they are becoming more extreme, as the dual-use nature of AI means many of its individual components most critical to national security are also commonplace in the commercial sector.

Meanwhile, U.S. regulatory capacity has not kept pace with technical developments, as the Departments of Commerce, Treasury, and State all lack sufficient technical and analytical capacity to effectively design and efficiently enforce technology protection policies on dual-use emerging technologies. Congress has taken some important steps in recent years to adapt technology protection regimes to challenges posed by emerging technologies, most notably the Export Control Reform Act of 2018 (ECRA) and the Foreign Investment Risk Review Modernization Act of 2018 (FIRRMA).⁷ However, more than two years after their passage, implementation of key aspects of both laws remains unfinished, hindering enforcement and confounding the affected industries.⁸

These conditions present policymakers with a difficult choice between under-protection, which will give competitors unacceptable levels of access to sensitive technologies, and over-protection, which has the potential to stifle innovation and harm overall U.S. competitiveness. Effective controls must target choke points that impose significant trickle-down strategic costs on competitors but minimal economic costs on U.S. industry. But such choke points are increasingly elusive.

Clearly state overarching principles to guide future U.S. dual-use technology protection policies. The United States must take a smarter and more predictable approach to applying technology protection policies to AI. The government should state that future technology protection policies will be guided by four overarching principles:

Recommendation

- U.S. technology controls must not supplant investment and innovation.
- U.S. strategies to promote and protect U.S. technology leadership must be integrated.
- The United States must be judicious in applying export controls to AI-related technologies, targeting discrete choke points and coordinating policies with allies.
- The United States must broaden investment screening on AI-related technologies.

“The United States must be judicious in applying export controls to AI-related technologies, targeting discrete chokepoints and coordinating policies with allies.”

On a technical level, AI poses particular challenges to control regimes given its dual-use, widespread, and largely open-source nature. Moreover, it builds on a host of other technologies. Given the ubiquitous nature of AI, export controls on AI algorithms carry substantial risk—improperly defined controls could inadvertently restrict the export of significant numbers of commercial products and cause substantial harm to the U.S. technology industry. While some AI algorithms are clearly candidates for export controls, such as those meant for use in battlefield applications, such software is largely already controlled under the Commerce Control List.⁹ Data is also a potential

target for controls—especially sensitive bulk data—but effective data controls face similar challenges to AI algorithms.¹⁰ Looking across the AI stack, the hardware component of the AI stack contains the most viable targets for traditional export controls.

Recommendation

Build regulatory capacity and fully implement ECRA and FIRRMA. The United States must also take steps to improve its capacity to design and implement effective technology protection policies. In the near term, the Departments of Commerce, Treasury, and State must ensure they have sufficient quantities of technically proficient personnel focused on technology protection policies and better utilize external advisory boards staffed with technical experts in designing policies. The Department of Commerce must also finalize its initial list of “emerging” and “foundational” technologies that must be controlled, as mandated by ECRA more than two years ago, and work to comprehensively adapt U.S. export control lists to address modern technology-focused security challenges.¹¹ Doing so is a critical step to implementing both ECRA and FIRRMA. Finally, departments and agencies should consider efforts to expedite and automate export licensing and Committee on Foreign Investment in the United States (CFIUS) filing proceedings, which could improve the effectiveness and reduce the economic costs of these regimes.¹²

Recommendation

Require investors from U.S. competitors to disclose transactions in a broader set of sensitive technologies to CFIUS. The United States must amend CFIUS’ authorities and procedures to enable it to better address modern challenges associated with sensitive, dual-use technologies. Specifically, it must enhance its ability to monitor investments from competitors in critical U.S. technology industries to prevent theft of IP and ensure that the United States retains control of sensitive technologies. U.S. competitors are investing heavily in U.S. AI firms. From 2010 to 2017, China-based investors poured more than \$1.3 billion into U.S. AI startups, and AI remains among the top technology areas for venture capital investment in the United States by China-based firms.¹³ However, the U.S. government has limited insight into these transactions. CFIUS is responsible for screening foreign investments for national security risks, but it only requires firms to disclose investments when the U.S. firm produces an export-controlled good—which very few AI firms do.¹⁴ Therefore, many firms based in the U.S. competitor countries that invest in U.S. AI companies have no obligation to report their investments to CFIUS. While CFIUS has broad authority to unwind such transactions, it currently has no visibility before they are consummated—creating a significant technology transfer risk.

CFIUS should increase disclosure requirements for investments in sensitive technologies by firms from China and Russia. Congress should mandate that all investments originating from “countries of special concern,” to include China and Russia, in national security–relevant applications of AI and other “sensitive technologies” as defined by CFIUS, must be disclosed to allow CFIUS the opportunity to review them prior to the completion of the transaction. This list of sensitive technologies should be distinct and broader than the list of “emerging” and “foundational” technologies required under ECRA and include industries key to U.S. national security that face persistent threats from adversarial capital, specifically

national security–relevant applications of AI, semiconductors, telecommunications equipment, quantum computing, and biotechnology, as well as other sectors identified in Made in China 2025. De-linking investment screening from export controls acknowledges that these two tools can and should be applied in different ways, permitting more expansive investment screening while maintaining targeted export controls focused on choke points. Limiting the scope of the mandatory filing requirements for this broader set of technologies to firms only from select U.S. competitors would prevent over-regulation and preserve the free flow of capital, increase insight into China’s and Russia’s investments in critical technologies, deter state-sponsored IP theft, and preserve U.S. leadership in AI for national security purposes.¹⁵



“CFIUS should increase disclosure requirements for investments in sensitive technologies for firms from China and Russia.”

Utilize targeted export controls on key semiconductor manufacturing equipment (SME). Where possible, the United States should use export controls to prevent competitors from obtaining AI capabilities that would grant them strategic or military advantages. The primary U.S. export control target to constrain competitors’ AI capabilities should be sophisticated SME necessary to manufacture high-end chips. SME is a critical choke point and an attractive target for export controls for the following reasons:

Recommendation

- Advanced AI is increasingly dependent on high–end computing capabilities¹⁶;
- China relies on international firms for its supply of high–end semiconductors; and
- SME manufacturing is specialized and dominated by the United States and its allies.




“The primary U.S. export control target to constrain competitors’ AI capabilities should be sophisticated semiconductor manufacturing equipment (SME) necessary to manufacture high-end chips.”

China is the only U.S. competitor attempting to cultivate a domestic, cutting-edge microelectronics fabrication industry capable of producing advanced chips at scale. Slowing the growth of China’s high-end semiconductor manufacturing ability would set back its attempts to build a cutting-edge microelectronics industry capable of fabricating chips most useful for advanced applications of AI for defense. Coupled with the efforts to promote U.S. semiconductor leadership outlined in Chapter 13 of this report, this will further the Commission’s proposed U.S. policy goal of remaining two generations ahead of China in cutting-edge microelectronics design and fabrication. However, controls on general-purpose semiconductors are unlikely to be effective given the larger number of countries capable of producing such chips. If implemented unilaterally, such controls could harm the U.S. semiconductor industry.

Recommendation

Align the export control policies of the United States, the Netherlands, and Japan regarding SME. The sophisticated photolithography tools needed to produce chips at the 16nm node and below, particularly extreme ultraviolet (EUV) and argon fluoride (ArF) immersion lithography tools, are the most complex and expensive type of SME. These tools are even more specialized than SME writ large, and the United States, the Netherlands, and Japan control the entire market.¹⁷ The Departments of State and Commerce should work with the governments of the Netherlands and Japan to align the export licensing processes of all three countries regarding high-end SME, particularly EUV and ArF immersion lithography equipment, toward a policy of presumptive denial of licenses for exports of such equipment to China. This would slow China’s efforts to domestically produce 7nm or 5nm chips at scale and constrain China’s semiconductor production capability of chips at any node at or below 16nm—which the Commission assesses to be most useful for advanced AI applications—by limiting the capability of Chinese firms to repair or replace existing equipment.¹⁸

Utilize targeted end-use export controls and reporting requirements to prevent use of high-end U.S. AI chips in human rights violations. The United States must take steps to prevent and deter U.S. firms from wittingly or unwittingly enabling uses of AI that violate human rights. List-based controls are ill-suited for this task given the commercial nature of most AI equipment, as the vast majority of its uses are legitimate. However, end-use and end-user export controls could be more effective. Although end-use controls are unlikely to prevent the transfer of strategic technologies to U.S. competitors determined to obtain them, they could prevent or deter U.S. firms from allowing certain key pieces of equipment, particularly high-end chips, to be utilized in malicious AI applications. Reporting revealing that U.S.-made chips are powering a supercomputer in Xinjiang, China, used for mass surveillance of Uyghur populations and that firms in China have filed patents for facial recognition specifically targeting Uyghurs illustrates the need to more closely monitor how high-end U.S. enabling hardware is utilized.¹⁹



“The United States must take steps to prevent and deter U.S. firms from wittingly or unwittingly enabling uses of AI which violate human rights.”

The Department of Commerce should prohibit the export of specific, high-performing AI chips for use in mass-surveillance applications, compel U.S. firms that export such chips to certify that the buyer will not utilize them to facilitate human rights abuses, and require that firms submit quarterly reports to Commerce listing all such chip sales to China. This would not constitute a licensing requirement that would introduce uncertainty and cause delays, but rather a self-certification and semi-regular report from industry. Such an action would demonstrate the U.S. commitment to ethical and responsible uses of AI, promote ethical behavior among U.S. firms, and make it harder for bad actors to utilize advanced U.S. chips for nefarious purposes.²⁰



“China’s campaign to exploit U.S.-based research violates the research community’s core principles of integrity, openness, accountability, and fairness.”

Strengthening Research Protection.

The U.S. research enterprise should be protected as a national asset. China’s campaign to exploit U.S.-based research violates the research community’s core principles of integrity, openness, accountability, and fairness.²¹ U.S. response measures to counter the actions of China’s government are nascent.²² There is a need for more technically versed intelligence collection and analysis on threats to the science and technology sector and a need to disseminate that information more broadly.²³ Government agencies, law enforcement, and research institutions need ready access to tools and resources to conduct nuanced risk assessment and build transparency around specific threats and tactics. The government and research institutions share responsibility for protecting core values and countering malicious activities. Responses should be coordinated with like-minded allies and partners to reinforce norms around openness of fundamental research, research integrity, and protection of IP.

Strengthening the integrity of the research process will support the foundations of open research. However, if not approached thoughtfully, U.S. policy actions to counter technology transfer could harm U.S. competitiveness and global scientific progress. Countering the CCP’s actions does not require severing most ties between research communities in China and the United States. The United States benefits from collaboration by staying connected with cutting-edge work in China and welcoming their PhD-level top talent²⁴ that comes to study at U.S. universities and remains in the United States at rates of 85% to 90% after graduating.²⁵

Recommendation

Build capacity to protect the integrity of the U.S. research environment. Congress should start by passing the Academic Research Protection Act (ARPA) and establishing a government-sponsored center of excellence on research security.²⁶ The ARPA legislation would create a dedicated National Commission on Research Protection, improve dissemination of open-source intelligence relating to foreign threats, and facilitate the sharing of studies and practices between government and research organizations.

Coordinate research protection efforts internationally with allies and partners. China's efforts to acquire foreign technology extend far beyond the United States.²⁷ The Office of Science and Technology Policy, the Department of State, and the Department of Justice should coordinate with allies and partners to further information-sharing on detrimental academic collaboration with entities affiliated with China's People's Liberation Army (PLA) and develop multilateral responses to mitigate the harm from these actions. Such diplomacy should seek to reinforce global norms around commitment to open fundamental research, as formalized in the United States in National Security Decision Directive 189.²⁸ The United States should strive to build a coalition committed to this principle and to research integrity, sidelining those who do not abide by the values that underpin innovation and global science cooperation.

“The United States should strive to build a coalition committed to this principle and to research integrity, sidelining those who do not abide by the values that underpin innovation and global science cooperation.”

Bolster cybersecurity support for research institutions. Protection of research data and IP from cyber-enabled theft is perhaps the most important measure and the most easily achieved layer of security. This is particularly salient for AI, when theft of training data or trained models essentially provides access to a final product. Federal grant-making agencies should ease the ability of research institutions to maintain a baseline level of cybersecurity by issuing clear guidance, establishing incentives, and sharing state-of-the-art best practices and resources.


Agencies such as the Department of Homeland Security (DHS) and FBI should increase support to information-sharing constructs and provide timely and actionable alerts on cyber threats and intrusions.²⁹ In addition, the government should broker commercial cloud credits for universities to support secure data storage for research groups and laboratories conducting research known to be of high interest to foreign adversaries.

Recommendation

Counter foreign talent recruitment programs. China's national plan for AI development directs use of foreign talent recruitment programs as a means to create a "high ground" for China's AI experts.³⁰ These problematic programs have received increasing attention in recent years. Rather than offering legitimate competition for scientific talent through attractive job offers, many are constructed in a manner that contradicts U.S. norms of research integrity, violates rules around disclosure, and creates vectors for technology transfer.³¹

The programs often employ a model of "part-time" recruitment, in which participants retain positions in the United States while accepting an affiliation with an institution in China.³² This often involves signing contracts that create conflicts through requirements to attribute patents to an institution in China, even if the research was conducted with U.S. funding. Participants often train other talent recruitment program members and replicate U.S.-funded work at an institution in China.

We commend recent action by Congress to limit the detrimental impact of these programs by mandating standardized disclosure requirements for federally funded research that will require comprehensive disclosure of conflicts of interest, conflicts of commitment, and all outside and foreign support.³³ This should be strengthened by a standardization and unification of grant application and documentation processes in machine-readable formats. Together, these measures would enable effective oversight, automated fraud



“Disclosure and grant standardization should be complemented with mandated and resourced compliance operations at each research funding agency—creating a layer of accountability to enforce disclosure policies and deter bad actors.”

detection, and data sharing across the federal research funding agencies. Disclosure and grant standardization should be complemented with mandated and resourced compliance operations at each research funding agency—creating a layer of accountability to enforce disclosure policies and deter bad actors.

Strengthen visa vetting to limit problematic research collaborations. Some U.S. universities and researchers are unknowingly entering into collaborative research arrangements with researchers from universities in China with close ties to the PLA and conducting research that directly contributes to China's military and security capabilities.³⁴ China's military-civilian fusion strategy and pursuit of technological leadership has been supported by a push from PLA-affiliated research institutions to send personnel abroad. Visiting scholars or students have been found to downplay ties to the military or deliberately obscure affiliation by using alternate names for their home institutions.³⁵

The United States should guard against the entrance of researchers with problematic affiliations through implementation of a special review process for visa applications from advanced-degree students and researchers with ties to research institutions affiliated with foreign military and intelligence organizations of designated countries of concern.³⁶ This should be accompanied by adequate resources to enable heightened review and paired with penalties that ban entry to visa applicants found to have intentionally not disclosed or improperly disclosed their military and intelligence affiliations.

Recommendation

Chapter 14 - Endnotes

¹ *Outline of the National Innovation-Driven Development Strategy*, Central Committee of the Communist Party of China and the PRC State Council (May 19, 2016) (translation by CSET on Dec. 11, 2019), <https://cset.georgetown.edu/research/outline-of-the-national-innovation-driven-development-strategy/>.

² *Deputy Assistant Attorney General Adam S. Hickey of the National Security Division Delivers Remarks at the Fifth National Conference on CFIUS and Team Telecom*, U.S. Department of Justice (April 24, 2019), <https://www.justice.gov/opa/speech/deputy-assistant-attorney-general-adam-s-hickey-national-security-division-delivers-0>.

³ *China Theft of Technology is Biggest Law Enforcement Threat to U.S., FBI Says*, The Guardian (Feb. 6, 2020), <https://www.theguardian.com/world/2020/feb/06/china-technology-theft-fbi-biggest-threat> (quoting William Evanina, director of the National Counterintelligence and Security Center).

⁴ A recent JASON study on the issue found that “[a]ctions of the Chinese government and its institutions that are not in accord with U.S. values of science ethics have raised concerns about foreign influence in the U.S. academic sector . . . there are problems with respect to research transparency, lack of reciprocity in collaborations and consortia, and reporting of commitments and potential conflicts of interest, related to these actions.” JASON, *Fundamental Research Security*, MITRE Corporation at 39 (Dec. 2019), https://www.nsf.gov/news/special_reports/jasonse13curity/JSR-19-2IFundamentalResearchSecurity_12062019FINAL.pdf.

⁵ In 2018 the U.S. National Counterintelligence and Security Center stated: “The threat to U.S. technology from Russia will continue over the coming years as Moscow attempts to bolster an economy struggling with endemic corruption, state control, and a loss of talent departing for jobs abroad.” *Foreign Economic Espionage in Cyberspace*, National Counterintelligence and Security Center at 8 (2018), https://s3-us-west-2.amazonaws.com/cyberscoop-media/wp-content/uploads/2018/07/26114025/2018ForeignEconomic-Espionage-Pub_FINAL.pdf.

⁶ Robert Farley & J. Tyler Lovell, *China’s Air Force Is Being Held Back by Its Terrible Jet Engines*, The National Interest (April 3, 2020) <https://nationalinterest.org/blog/buzz/chinas-air-force-being-held-back-its-terrible-jet-engines-140252>.

⁷ See Pub. L. 115-232, Title XVII, Subtitle B, 132 Stat. 1636, 2208, as amended by Pub. L. 116-6, Division H, Title II, Section 205 Consolidated Appropriations Act, 2019, 133 Stat. 13, 476; Pub. L. 115-232, Title XVII, Subtitle A, 132 Stat. 1636, 2174.

⁸ Of particular note, ECRA requires the Department of Commerce to identify “emerging and foundational technologies that are essential to the national security of the United States” that are not otherwise controlled, but to date Commerce has not identified a single technology under this provision. This has left gaps in the U.S. approach to protecting its advantages in high-tech sectors, including AI, and created uncertainty for industry. See Pub. L. 115-232, Title XVII, Subtitle B, 132 Stat. 1636, 2208, as amended by Pub. L. 116-6, Division H, Title II, Section 205 Consolidated Appropriations Act, 2019, 133 Stat. 13, 476.

⁹ Carrick Flynn, *Recommendations on Export Controls for Artificial Intelligence*, Center for Security and Emerging Technology (Feb. 6, 2020), <https://cset.georgetown.edu/research/recommendations-on-export-controls-for-artificial-intelligence/>.

¹⁰ There is also room to work with allies and partners to create standards for securely transferring key data sets and limiting their distribution to certain trusted nations; see Chapter 15 of this report for additional details on this topic.

¹¹ See Chris Darby, et al., *Mitigating Economic Impacts of the COVID-19 Pandemic and Preserving U.S. Strategic Competitiveness in Artificial Intelligence*, NSCAI at 16 (May 19, 2020), <https://www.nsc.gov/white-papers/covid-19-white-papers/>.

¹² The Chapter 14 Blueprint for Action contains more details on this recommendation.

¹³ Chinese venture capital investment in the U.S. increased substantially after 2014 but has stalled since 2018. Nevertheless, AI remains one of the top sectors for Chinese venture capital investment in the U.S. Michael Brown & Pavneet Singh, *China's Technology Transfer Strategy*, Defense Innovation Unit Experimental (Jan. 2018), [https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_\(1\).pdf](https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_(1).pdf); Adam Lysenko, et al., *Disruption: US-China Venture Capital in a New Era of Strategic Competition*, (Jan. 2020), https://publications-research.s3-us-west-2.amazonaws.com/RHG_Disruption_US+China+VC_January2020.pdf; Mercedes Ruehl, et al., *Chinese State-Backed Funds Invest in U.S. Tech Despite Washington Curbs*, Financial Times (Dec. 2, 2020), <https://www.ft.com/content/745abeca-561d-484d-acd9-ad1caedf9e9e>.

¹⁴ As a result, CFIUS disclosure requirements disproportionately impact investments from U.S. allies; of the 94 mandatory CFIUS filings in 2019, 14 were from Japan, 12 were from Canada, 11 were from the U.K., and only three were from China. *Annual Report to Congress, CFIUS at 33-36* (2019), <https://home.treasury.gov/system/files/206/CFIUS-Public-Annual-Report-CY-2019.pdf>.

¹⁵ This would require an amendment to CFIUS' authorizing legislation, the draft text of which can be found in the Legislative Text Appendix of this report. Specifically, it would amend Section 721(a) of the Defense Production Act of 1950 (as amended and codified in 50 USC § 4565[a]) and grant the Department of the Treasury new authorities to alter mandatory filing requirements and define a list of "sensitive technologies" distinct from export control lists.

¹⁶ OpenAI estimates that since 2012, the amount of compute used in the largest AI training runs is doubling every 3.4 months. See Dario Amodei & Danny Hernandez, *AI and Compute*, OpenAI (May 16, 2018), <https://openai.com/blog/ai-and-compute/>.

¹⁷ The Dutch firm ASML has a monopoly on EUV lithography tools, which are the most advanced type, and ArF immersion lithography tools are only produced by ASML and the Japanese firm Nikon.

¹⁸ The Wassenaar Arrangement lists lithography equipment capable of making chips with features of 45nm or below as a controlled item. However, because the Wassenaar Arrangement is not binding, states parties are not obligated to comply with this as a legal restriction. See *List of Dual-Use Goods and Technologies and Munitions List*, Wassenaar Arrangement Secretariat at 72 (Dec. 2018), <https://www.wassenaar.org/app/uploads/2019/consolidated/WA-DOC-18-PUB-001-Public-Docs-Vol-II-2018-List-of-DU-Goods-and-Technologies-and-Munitions-List-Dec-18-1.pdf>.

¹⁹ Paul Mozur & Don Clark, *China's Surveillance State Sucks Up Data. U.S. Tech Is Key to Sorting It*, New York Times (Nov. 24, 2020), <https://www.nytimes.com/2020/11/22/technology/china-intel-nvidia-xinjiang.html>; Leo Kelion, *Huawei Patent Mentions Use of Uighur-spotting Tech*, BBC (Jan. 13, 2021), <https://www.bbc.com/news/technology-55634388>. Reporting indicates firms in China may be in the process of altering these patents to remove references to specific ethnic groups.

²⁰ This action would build on recent State Department guidance regarding best practices for transactions linked to foreign government end-users for products or services with surveillance capabilities. See *U.S. Department of State Guidance on Implementing the "UN Guiding Principles" for Transactions Linked to Foreign Government End-Users for Products or Services with Surveillance Capabilities*, U.S. Department of State (Sept. 30, 2020), <https://www.state.gov/key-topics-bureau-of-democracy-human-rights-and-labor/due-diligence-guidance/>.

²¹ JASON, *Fundamental Research Security*, MITRE Corporation (Dec. 2019), https://www.nsf.gov/news/special_reports/jasonsecurity/JSR-19-2IFundamentalResearchSecurity_12062019FINAL.pdf.

²² Promising efforts have been initiated through the National Counterintelligence Task Force and the Office of Science and Technology Policy's Joint Committee on Research Environments, as well as among universities, to build communities of interest to share best practices and conduct internal audits around disclosure policies and cybersecurity. See National Science and Technology Council, The White House (last accessed Jan. 1, 2021), <https://www.whitehouse.gov/ostp/nstc/>, and the Academic Security and Counter Exploitation Program launched by the Texas A&M University System, *Academic Security and Counter Exploitation Program*, Texas A&M University (last accessed Jan. 11, 2021), <https://asce.tamus.edu/>.

²³ Chapter 5 of this report recommends that the Intelligence Community (IC) should prioritize and accelerate collection of scientific and technical intelligence to better understand adversary capabilities and intentions.

Chapter 14 - Endnotes

²⁴ The Commission supports measures to strengthen the ability of the United States to attract and retain top AI talent coming from China and elsewhere. See Chapter 10 of this report.

²⁵ Remco Zwetsloot, *U.S.-China STEM Talent "Decoupling,"* Johns Hopkins Applied Physics Laboratory at 13 (2020), <https://www.jhuapl.edu/assessing-us-china-technology-connections/dist/407b0211ec49299608551326041488d4.pdf>.

²⁶ See H.R. 8346, Academic Research Protection Act, 116th Congress (2020), <https://www.congress.gov/bill/116th-congress/house-bill/8346>. The legislation sought to establish a National Commission on Research Protection; establish an open-source intelligence clearinghouse relating to foreign threats to academic research overseen by the Director of National Intelligence (DNI); improve guidance from the Departments of State and Commerce to ensure academic institutions are meeting export-control responsibilities; and develop a Federal Bureau of Investigation (FBI) outreach strategy on threats to the academic community.

²⁷ Notably, two-thirds of overseas professional associations that transfer technology to China are located outside the United States, mainly distributed among U.S. allies and partners. Ryan Fedasiuk & Emily Weinstein, *Overseas Professionals and Technology Transfer to China*, Center for Security and Emerging Technology at 11 (July 2020), <https://cset.georgetown.edu/research/overseas-professionals-and-technology-transfer-to-china/>.

²⁸ The directive defines fundamental research as "Basic and applied research in science and engineering, the results of which ordinarily are published and shared broadly within the scientific community, as distinguished from proprietary research and from industrial development, design, production, and product utilization, the results of which ordinarily are restricted for proprietary or national security reasons." *National Policy on the Transfer of Scientific, Technical, and Engineering Information*, Executive Office of the President (Sept. 21, 1985), <https://fas.org/irp/offdocs/nsdd/nsdd-189.htm>.

²⁹ Such as the Research and Education Networks Information and Sharing Analysis Center (REN-ISAC). REN-ISAC (last accessed Jan. 2, 2021), <https://www.ren-isac.net/>.

³⁰ William C. Hannas & Huey-meei Chang, *China's Access to Foreign AI Technology*, Center for Security and Emerging Technology at 9-10 (Sept. 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_China_Access_To_Foreign_AI_Technology.pdf.

³¹ The Office of Science and Technology Policy defines foreign government talent recruitment programs as “an effort directly or indirectly organized, managed, or funded by a foreign government to recruit science and technology professionals or students (regardless of citizenship or national origin).” *Enhancing the Security and Integrity of America’s Research Enterprise*, White House Office of Science and Technology Policy at 18 (July 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/07/Enhancing-the-Security-and-Integrity-of-Americas-Research-Enterprise.pdf>.

³² David Zweig & Siqin Kang, *America Challenges China’s National Talent Programs*, Center for Strategic and International Studies at 5 (May 2020), https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/20505_zweig_AmericaChallenges_v6_FINAL.pdf?bTLm4WdtG93lAVmxLdlWsgkgeNQDQUAv.


³³ See Pub. L. 116-283, sec. 223, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

³⁴ Glen Tiffert, *Global Engagement: Rethinking Risk in the Research Enterprise*, The Hoover Institution (July 2020), https://www.hoover.org/sites/default/files/research/docs/tiffert_globalengagement_full_0818.pdf. A subsequent study of a larger database of research papers conducted by Nature identified more than 12,000 papers from the years 2015 to 2019 that were co-authored by researchers in the U.S. and at one of China’s “Seven Sons” universities. Furthermore, *Nature* found that “among those, 499 authors had a dual affiliation with a U.S. institution and a Seven Sons university and were listed on papers declaring grant funding from the NIH or the U.S. National Science Foundation.” Nidhi Subbaraman, *US Investigations of Chinese Scientists Expand Focus to Military Ties*, *Nature* (Sept. 9, 2020), <https://www.nature.com/articles/d41586-020-02515-x>.

³⁵ Alex Joske, *Picking Flowers, Making Honey: The Chinese Military’s Collaboration with Foreign Universities*, Australian Strategic Policy Institute (Oct. 2018), <https://www.aspi.org.au/report/picking-flowers-making-honey>.

³⁶ The Commission recommends this as an update to Presidential Proclamation 10043 that suspends F or J visas to study or conduct research for Chinese nationals affiliated with the Chinese government military-civil fusion strategy. 85 Fed. Reg. 34353, *Suspension of Entry as Nonimmigrants of Certain Students and Researchers from the People’s Republic of China*, Executive Office of the President (May 29, 2020), <https://www.federalregister.gov/documents/2020/06/04/2020-12217/suspension-of-entry-as-nonimmigrants-of-certain-students-and-researchers-from-the-peoples-republic>.

Chapter 15: A Favorable International Technology Order



Coordination and Alignment with Allies and Partners

Develop an International S&T Strategy



Reorient the Department of State and U.S. Foreign Policy



Launch an International Digital Democracy Initiative



Build an Emerging Technology Coalition



Implement a Comprehensive U.S. National Plan



Cultivate an International Emerging Technology Research Hub



The United States cannot compete with and counter the global technology ambitions of authoritarian rivals if it acts alone. Like-minded countries must work together to advance an international rules-based order, protect free and open societies, and unleash economic innovation. The authoritarian challenge to the global technology order encompasses five distinct but related elements:

- A rising challenge to U.S. and Western technology firms for global market share, impacting the prosperity and global economic position of the United States and its allies and partners;
- China's increasing influence and strategic leverage over countries that utilize technologies and infrastructure built and developed in China;
- The prospect of authoritarian consolidation in states that gain easy access to digital tools that can strengthen repressive rule;
- The prospect of democratic backsliding in states with governments that may be tempted to utilize digital tools in ways that undermine liberal values; and
- A threat to the cohesion of democratic allies as an influential bloc of states with the capacity to shape global technology norms and standards.¹



“The United States must pursue a comprehensive strategy in close coordination with our allies and partners for AI innovation and adoption that promotes values critical to free and open societies.”

The United States must pursue a comprehensive strategy in close coordination with our allies and partners for artificial intelligence (AI) innovation and adoption that promotes values critical to free and open societies. Furthermore, the United States must collaborate with its closest allies and partners to develop principles for employing AI tools ethically and responsibly, defend the integrity of international technical standards, promote digital markets, leverage comparative expertise to develop privacy-preserving technologies, and share practices and resources to defend against authoritarian attacks on digital infrastructure and democratic values.

To achieve these goals, the Commission proposes that the White House request the Department of State to lead an effort with and other key agencies to:

Develop and implement an International Science and Technology Strategy (ISTS) to help coordinate AI and emerging-technology policies government-wide and with our closest allies and partners; apply the tools of foreign assistance, technical expertise and guidance, and development finance; and foster collaborative R&D. The ISTS should serve as the international component of the National Technology Strategy (see Chapter 9 of this report). The ISTS should be centered around four big initiatives:

Recommendation

- **Build an Emerging Technology Coalition** of allies and partners to promote the design, development, and use of emerging technologies according to democratic norms and values; coordinate policies and investments to counter the malign use of these technologies by authoritarian regimes; and provide concrete, competitive alternatives to counter the adoption of digital infrastructure made in China.
- As part of the Emerging Technology Coalition, **launch an International Digital Democracy Initiative** with allies and partners to align international assistance efforts to develop, promote, and fund the adoption of AI and associated technologies that comports with democratic values and ethical norms around openness, privacy, security, and reliability.
- **Implement a comprehensive U.S. national plan to support international technology efforts** around technical standards, foreign assistance, development finance, and export controls.
- **Enhance the United States' position as an international emerging technology research hub** for collaborative R&D efforts by formalizing a partnership between the U.S. National AI Research Institutes and multilateral initiatives like the Global Partnership on AI (GPAI), creating a Multilateral AI Research Institute (MAIRI) in the United States with key allies and partners, and catalyzing international collaboration and talent exchanges.

Build an Emerging Technology Coalition. The United States should lead an Emerging Technology Coalition (ETC) of like-minded nations either as part of a larger democracy summit or as a stand-alone endeavor. The immediate step for the ETC should be to organize its efforts to synchronize policies around the following seven critical areas:


Recommendation

- **Developing and operationalizing standards and norms**, in support of democratic values and the development of secure, reliable, and trusted technologies;
- **Promoting and facilitating coordinated and joint R&D on AI and digital infrastructure** that advances shared interests and benefits humanity;
- **Promoting democracy, human rights, and the rule of law** through joint efforts to counter censorship, malign information operations, human trafficking, and illiberal uses of surveillance technologies;
- **Exploring ways to facilitate data-sharing** among allies and partners through enabling agreements, common data archival procedures, cooperative investments in privacy-enhancing technologies, and addressing legal and regulatory barriers;
- **Promoting and protecting innovation**, particularly through export controls, investment screening, supply chain assurance, emerging technology investment, trade policy, research and cyber protections, and intellectual property alignment;
- **Developing AI-related talent**, by analyzing labor market challenges, harmonizing skills and certification requirements, and increasing talent exchanges, joint training, and workforce development initiatives;
- **Launching the International Digital Democracy Initiative.**²

Recommendation

Launch an International Digital Democracy Initiative (IDDI). As part of the ETC, the United States, with its allies and partners, should launch an IDDI to align international assistance efforts to develop, promote, and fund the adoption of AI and associated technologies that comport with democratic values and ethical norms around openness, privacy, security, and reliability.

IDDI Strategy.



The IDDI will be critical for enabling nations around the world to adopt secure, trusted, and open digital ecosystems,³ empowering communities to use AI and digital technologies in ways that strengthen democracies, promote sustainable development, and advance shared values like privacy, human rights, and the rule of law. IDDI further provides an opportunity for the United States and like-minded allies and partners to counter authoritarian uses of AI, particularly by providing alternatives to digital infrastructure projects that are used for illiberal ends, endanger the social cohesion among and between democracies, and threaten collective security.⁴ As international digital and telecommunications infrastructure investment needs continue to grow⁵ and China continues to use digital development to export authoritarianism and expand influence, the United States and its allies and partners must join forces to coordinate a strategy that maximizes the impact of government assistance efforts and also catalyzes private-sector investment to address shared challenges.

“... the United States and its allies and partners must join forces to coordinate a strategy that maximizes the impact of government assistance efforts and also catalyzes private sector investment to address shared challenges.”

Implement a comprehensive U.S. national plan to support international technology efforts. The ISTS should include an integrated government-wide plan for using and bolstering the tools of U.S. foreign policy—including technical and foreign assistance, development financing, and export controls—to advance the ETC, the IDDI, and stand-alone projects. As demonstrated below, the plan should include methods to shape international technical standards; coordinate and expand programs of the Department of State, the United States Agency for International Development, the U.S. International Development Finance Corporation, and other federal agencies; and use targeted export controls to preserve key U.S. and allied technical advantages and also further transparency and accountability. It will require significant, dedicated appropriations to achieve meaningful results.⁶

Recommendation



International Science & Technology Strategy (ISTS) Task Force

Convened by the White House with leadership from State, Treasury, Commerce, Energy, DFC, EXIM, MCC, NSF, USAID, USTDA, and other critical agencies

Develop and oversee implementation of U.S. government-wide strategy for international technical standards and international digital development efforts

Select U.S. Stakeholders and Proposed Roles



Department of State

Foreign policy leadership and diplomacy

Senior leadership and Ambassador-at-Large lead efforts to establish and implement ETC and IDDI

Implement holistic effort to coordinate international security, economic policy, S&T, human rights, foreign assistance

Facilitate tech diplomacy through U.S. Embassies and Missions



United States Agency for International Development (USAID)

Digital development and humanitarian assistance

Prioritize implementation of digital development through Digital Strategy

Provide resources, tools, and expertise for broader U.S. digital development projects

Advise on international technical standardization



Department of Commerce



National Institute of Standards and Technology (NIST)

Advise on international technical standardization

Coordinate interagency task force on technical standards

Improve partnership and collaboration with industry

Align standards for secure, reliable, and trusted technologies with key allies and partners



U.S. International Development Finance Corporation (DFC)

Foreign direct investment

Expand investments in technology and digital infrastructure

Increase blended finance transactions to achieve scale



U.S. Export-Import Bank (EXIM)

Export promotion and financing assistance

Leverage Program on China and Transformational Exports to strengthen U.S. tech competitiveness

Advise on incentivizing and exporting democratic emerging technologies



Bureau of Industry and Security (BIS)

End-user controls

Develop and coordinate end-user licensing policies and export controls to further democratic values



U.S. Trade and Development Agency (USTDA)

Export promotion and technical assistance

Support export of U.S. emerging technologies through increased funding, training, assistance, and pilot projects

Recommendation

Enhance the United States' position as an international emerging technology research hub.

The United States must maintain its leadership in international R&D by further establishing itself as a hub of international research into and involving emerging technologies to foster AI collaboration and coordination with key allies and partners. These efforts will facilitate critical support to the ETC and IDDI by developing digital technologies and best practices that comport with democratic values; enhance U.S. contributions to existing and future international efforts like GPAI; and provide avenues for the United States and allies—particularly European allies—to pool resources to address commercial gaps in R&D and



“The United States must maintain its leadership in international R&D by further establishing itself as a hub of international research into and involving emerging technologies to foster AI collaboration and coordination with key allies and partners.”

overcome challenges to collaboration around cross-border data-sharing. Making the United States an international emerging technology research hub has three components:



- **First, the United States should provide formal research support to key international efforts such as GPAI and the Organisation for Economic Co-operation and Development,**⁷ particularly through the National Science Foundation (NSF)’s National AI Research Institutes.⁸ The important research undertaken by the National AI Research Institutes—run by the NSF and other U.S. agencies—and by other United States departments and agencies is an incredible resource that should support these international efforts and advance AI and digital goals of the U.S. and like-minded partners.

- **Second, the United States should work with key allies and partners to establish the Multilateral AI Research Institute (MAIRI).** MAIRI will facilitate joint efforts to develop technologies that advance responsible, human-centric, and privacy-preserving AI/machine learning (ML) that better societies and allow allies to pool their talents and resources. It will provide a model for equitable, multilateral research, facilitate AI R&D that builds on like-minded countries' strengths, and foster a global AI workforce for the next generation. MAIRI will be key to a U.S.-led effort to promote values of free and open societies, win the global technology competition, unleash AI innovation and economic prosperity, and develop AI applications that benefit humanity. MAIRI members will champion agreed-upon research integrity principles, leverage trusted infrastructure and research resources, and seek to be a part of a federated network of global research institutes. NSF should be the anchor partner, but MAIRI should be structured to enable participation of other federal agencies, like the Departments of State and Energy.⁹ The United States should fund the initial startup costs, including acquisition of MAIRI's physical center located in the United States.
- **Third, the United States should leverage existing O and J visa programs to facilitate foreign researchers' involvement in joint projects.** Sustained, strong collaboration between the United States and allies and partners is critical to winning this technology competition and unleashing innovation and entrepreneurship across like-minded countries. There is no substitute for shoulder-to-shoulder research for building relationships, exchanging ideas and expertise, and sparking future collaboration.¹⁰


Recommendation

Reorient U.S. foreign policy and the Department of State for great power competition in the digital age. New outward-facing digital foreign policy initiatives are only part of the equation for ensuring the long-term success of global technology policy. The United States must make inward-focused reforms to the Department of State as well. There is currently no clear lead for emerging technology policy or diplomacy within the State Department, which hinders the Department's ability to make strategic technology policy decisions. It also creates confusion for allies and partners, who regularly express uncertainty regarding which senior official should be their primary point of contact for issues related to key topics such as AI, 5G, quantum computing, biotechnology, or new emerging technologies.

Competitive diplomacy in AI and emerging technology arenas is a strategic imperative in an era of great power competition and necessitates an intensified reorientation of the Department of State. The United States must redesign the internal structure, focus, and culture of the State Department to adapt U.S. diplomacy for the digital age and empower diplomats to advance American interests at the intersection of technology, security, commerce, and human rights. Supporting these efforts and succeeding in U.S. diplomacy will require targeted appropriations from Congress.

The Commission recommends the following immediate actions to reorient U.S. diplomacy:

- **First, the Deputy Secretary of State for Management and Resources (D/MR)** should have responsibility to prioritize reorienting and reorganizing the Department for technology diplomacy. Past administrations have used the D/MR position to lead on strategic priorities and ensure execution. The D/MR should provide direction for immediate and long-term planning around technology diplomacy, including policy development, coordination, and resourcing. The D/MR should also have a lead role in oversight and implementation of the ISTS.



“The United States must redesign the internal structure, focus, and culture of the State Department to adapt U.S. diplomacy for the digital age ...”

- **Second, the State Department should expedite and prioritize efforts to staff, resource, and build out the newly created Bureau of Cyberspace Security and Emerging Technologies (CSET Bureau).** The CSET Bureau, approved in early January 2021, would be led by an official with the title of Ambassador-at-Large and Coordinator. The bureau would have a critical role as the focal point for U.S. diplomatic efforts around security challenges associated with emerging technologies and would provide an accountable home for AI advocacy within the Department.¹¹ The Department, with congressional support, must ensure the CSET Bureau is adequately staffed and resourced. Quickly standing up CSET is critical to building the Department’s technology diplomacy capacity, improving technology policy coordination across the Department, and regularly elevating technology issues to the attention of senior leaders. The Department should assess where CSET should be placed to best achieve those objectives but must ensure its creation is not further delayed.
- **Third, the State Department should enhance its presence in foreign and U.S. technology hubs** with a cadre of dedicated technology officers at U.S. missions to strengthen diplomatic advocacy, improve technology scouting, and inform policy and foreign assistance.
- **Fourth, AI-related technology modules should be incorporated into Foreign Service Institute training courses** at multiple levels to ensure U.S. diplomats are equipped to lead in an environment being transformed by emerging technology.
- **Fifth, Congress must appropriate funds necessary for urgent State Department needs** both to augment the U.S. diplomatic corps and to support critical State Department programs focused on AI and emerging technologies to advance U.S. interests.

These steps are essential, but not sufficient, to further U.S. interests in tech diplomacy. Eventually, the D/MR role should transition into a permanent Under Secretary for Science, Research and Technology (State/Q). State/Q would lead a reorganization of the Department, combining offices and bureaus to further a robust, coordinated approach to science and technology diplomacy and foreign assistance in the context of great power competition.¹²

Chapter 15 - Endnotes

¹ The threat to allied cohesion also extends to the military realm, insofar as building divergent or incompatible digital systems poses challenges for interoperability or creates risks for U.S. forces operating in allied countries. See Daniel Kliman, *Why the United States Needs a Digital Development Fund*, Center for a New American Security at 2 (Oct. 10, 2019), <https://www.cnas.org/publications/commentary/why-the-united-states-needs-a-digital-development-fund> (“Over the long term, China’s digital investments could render some developing countries off-limits to U.S. forces, constricting the geography of American military access.”).

² Detail on these critical areas can be found in the Chapter 15 Blueprint for Action and its associated Annex.

³ USAID’s Digital Strategy defines the “digital ecosystem” as the “stakeholders, systems, and enabling environments that together empower people and communities to use digital technology to gain access to services, engage with each other, or pursue economic opportunities.” This includes “a sound enabling environment and policy commitment; robust and resilient digital infrastructure; capable digital service providers and workforce; and, ultimately, empowered end-users of digitally enabled services.” *Digital Strategy 2020-2024*, USAID at 4 (June 2020), <https://www.usaid.gov/usaid-digital-strategy>.

⁴ The Chinese government’s global infrastructure projects and its widespread state influence within its private sector have enabled Chinese firms to provide surveillance and smart-city technologies to hundreds of cities globally, particularly in developing countries, bolstering autocratic regimes and enabling Chinese geopolitical coercion and government data collection. See, e.g., Hugh Harsono, *China’s Surveillance Technology is Keeping Tabs on Populations Around the World*, *The Diplomat* (June 18, 2020), <https://thediplomat.com/2020/06/chinas-surveillance-technology-is-keeping-tabs-on-populations-around-the-world/>; Testimony of Steven Feldstein before the U.S.-China Economic and Security Review Commission, *Hearing on China’s Strategic Aims in Africa* (May 8, 2020), https://www.uscc.gov/sites/default/files/Feldstein_Testimony.pdf.

⁵ To support the G20, the Global Infrastructure Hub has forecasted global telecommunications infrastructure investment needs at \$8.9 trillion over the next approximately 20 years, with current trends falling short of these needs by \$1 trillion. *Forecasting Infrastructure Investment Needs and Gaps*, Global Infrastructure Hub (last accessed Jan. 13, 2021), <https://outlook.gihub.org/>.

⁶ Detailed recommendations for U.S. agencies and Congress can be found in the Chapter 15 Blueprint for Action.

⁷ GPAI was launched in 2020 to “foster responsible development of AI grounded in these principles of human rights, inclusion, diversity, innovation and economic growth.” Current members include Australia, Brazil, Canada, the European Union, France, Germany, India, Italy, Japan, Mexico, the Netherlands, New Zealand, Poland, Singapore, Slovenia, South Korea, Spain, the United Kingdom, and the United States, with the OECD and UNESCO as Permanent Observers. GPAI bridges “the gap between theory and practice,” particularly through research and technical expertise shared via multi-stakeholder working groups. *About GPAI*, GPAI (last accessed Jan. 6, 2020), <https://www.gpai.ai/about/>; *UNESCO Joins Global Partnership on Artificial Intelligence as Observer*, UNESCO (Dec. 10, 2020), <https://en.unesco.org/news/unesco-joins-global-partnership-artificial-intelligence-observer>.

⁸ *Artificial Intelligence at NSF*, NSF (Aug. 26, 2020), <https://www.nsf.gov/cise/ai.jsp>.

⁹ For example, the Department of Energy may provide critical expertise on undertaking applied research with industry or through its national laboratories, particularly on high-performance and quantum computing, while the Department of State can provide foreign policy expertise and support initiatives on data-sharing and AI research clouds with allies and partners.

¹⁰ Detailed recommendations for each of these components can be found in the Chapter 15 Blueprint for Action.

¹¹ *Secretary Pompeo Approves New Cyberspace Security and Emerging Technologies Bureau*, U.S. Department of State (Jan. 7, 2021), <https://2017-2021.state.gov/secretary-pompeo-approves-new-cyberspace-security-and-emerging-technologies-bureau//index.html>.

¹² These components of the State Department should include key functions of the CSET Bureau; the Bureau of Oceans, Environment and Science; the Office of the Science and Technology Adviser; the Coordinator for Cyber Issues; and the Center for Analytics.

Chapter 16: Associated Technologies

AI and Beyond: Identifying and Prioritizing Critical Emerging Technologies



Artificial Intelligence



Biotechnology



Robotics and Autonomy



Advanced Manufacturing



Semiconductors



Quantum Computing



5G and Advanced Networking



Energy Systems

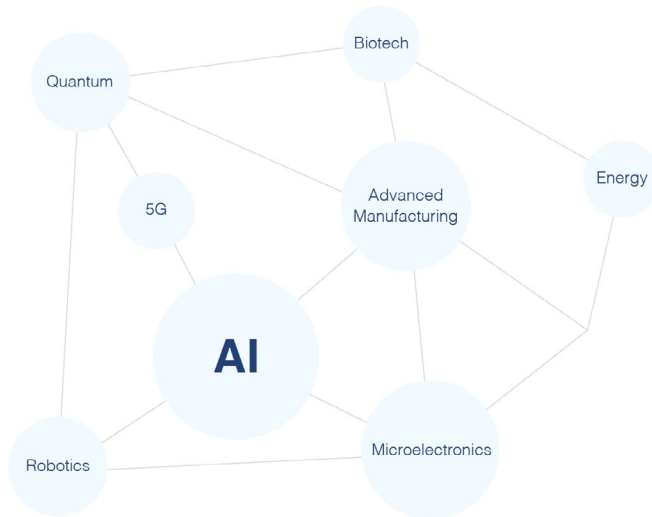
The Commission's work ends where it began, with the conclusion that artificial intelligence (AI) will transform virtually every aspect of our existence. However, leadership in AI is not an end unto itself; it is a necessary, but not sufficient, condition for the overarching goal of preserving U.S. leadership in technology. That reality presents a challenge for U.S. strategy: how to prioritize investments in AI and other key emerging technologies and support specific projects that will build on and amplify cross-technology strengths. The United States must view its efforts to lead in AI through the broader lens of competition across a range of emerging technologies, and, therefore, also support a comprehensive strategy to sustain U.S. leadership in key associated technologies.



“The United States must view its efforts to lead in AI through the broader lens of competition across a range of emerging technologies ...”

Leadership in AI relies on and drives leadership across a suite of emerging technologies. AI sits at the center of the constellation of emerging technologies, enabling some and being enabled by others.¹ For instance, 5G and quantum computing are poised to enable new growth in AI capabilities, while AI stands to transform the biological sciences, producing significant technological breakthroughs and turning the biotechnology sector into one of the primary drivers of overall economic competitiveness.²

AI is at the center of a constellation of emerging technologies, enabling some and enabled by others. The United States Government must view its efforts to lead in AI as part of a broader strategy to compete in each of these technologies:



The United States Government must:

Define and prioritize the key emerging technologies needed to ensure U.S. national competitiveness, including U.S. leadership in the following technologies and platforms:

- Advanced biofabrication capabilities
- Quantum chip fabrication
- 5G spectrum sharing
- Robotics software
- Additive manufacturing
- Energy storage technologies

Associated Technologies.

China is pursuing a comprehensive technology leadership strategy. China's strategic investments in key sectors through its Made in China 2025 initiative threaten U.S. technological prowess, economic prosperity, and national security.³ In addition to investments in AI, China is seeking to become a world leader in quantum, 5G, and biotech, among other areas, and sees its strategies to lead in AI and each of these other technologies as mutually reinforcing. It has made clear which technologies it views as key national priorities and is investing heavily in a wide range of sectors that it assesses are essential to overall technical leadership.

The United States has neither identified, nor prioritized leadership in, the technologies that are central to national competitiveness. The United States must develop a unified list of the technologies that will underpin national competitiveness in the 21st century because the first-mover advantage of developing and deploying technologies like microelectronics, biotechnology, and quantum computing will make it difficult for the United States to catch up to China. The lack of such a list results in disparate funding and policy approaches to technology protection and promotion across the U.S. government. The absence of clear priorities also makes it more difficult to effectively marshal private-sector investment in key technologies. In critical sectors with strong network effects like telecommunications, a winner-take-all dynamic raises the stakes even further for rapidly developing a leading technology platform.⁴

Ensuring U.S. leadership in the manufacturing of key emerging technology platforms will be an essential component of national competitiveness. Identifying and supporting research in priority technologies is a necessary but insufficient step to maintain national competitiveness in emerging technologies. The United States will also have to invest in the production of strategic physical elements of these technologies to create game-changing platforms, maximize U.S. competitiveness, and reduce dependencies that create national security vulnerabilities. Such investments are often expensive, but a strategic approach does not require manufacturing every advanced component domestically and will pay tremendous long-term dividends. The need to support advanced manufacturing applies to nearly every key emerging technology sector, including semiconductors, quantum computing, biotechnology, telecommunications equipment, and others, and is reflected in the following recommendations.


Technology leadership will require major new investments in underlying digital infrastructure. It is impossible to ignore the state of America's underlying digital infrastructure when considering a strategy to preserve overall U.S. leadership in technology. The sophistication and reach of the core U.S. digital infrastructure that underpins connectivity, namely high-speed internet and telecommunications networks, significantly lag behind those of many other developed nations.⁵ Additionally, U.S. physical infrastructure remains largely disconnected; no U.S. cities are ranked among the top 10 in smart-city connectedness, and only one is in the top 30.⁶ Maximizing citizens' access to the digital economy, ensuring they have the requisite digital skills, and more closely connecting the physical and digital worlds will be necessary to fuel future growth. Boosting the digital connectivity of physical U.S. assets will not only enhance their effectiveness and reliability, but also generate new sources of data that enable novel, potentially revolutionary uses of AI in areas ranging from energy grid management and urban planning to transportation. As the United States considers options to modernize U.S. physical infrastructure, prioritizing its digital connectivity will provide substantial long-term benefits and buttress U.S. technology competitiveness and national security.

Identifying and Prioritizing Technologies Central to National Competitiveness.

While the United States should by no means adopt China's centrally planned and state-directed economic model, it must start by developing better strategic planning, forecasting, and prioritization of emerging technologies to ensure long-term competitiveness. The government should:

Recommendation

Define and prioritize the key emerging technologies that are needed to ensure U.S. national competitiveness. As part of its National Technology Strategy (see Chapter 9 of this report), the White House should publish a list of critical and emerging technologies in which U.S. leadership is essential. It should develop detailed implementation plans for each sector to determine how the government should best work with industry to promote U.S. leadership, assess which specific subsectors are crucial to national security, and determine what



“As part of its National Technology Strategy (see Chapter 9), the White House should publish a list of critical and emerging technologies in which U.S. leadership is essential.”

regulatory steps or incentives are necessary to create the required investment environment. These plans should promote investment in specific platforms that will have a force multiplier effect on U.S. technology leadership, identify key choke points where competitors could potentially be blocked with minimal impact on U.S. industry, and promote supply chain resiliency. The creation and maintenance of such a list and the associated implementation plans will help produce a national consensus across government, industry, and academia about which sectors are most important in the emerging techno-economic competition. The result will be an important message to Congress regarding where the country must prioritize and expend resources, as well as a powerful demand signal to industry.

Many similar lists exist throughout the government, but there has been no effort to unify them into a single, authoritative document accompanied by a strategic vision and detailed follow-through actions designed to ensure long-term U.S. leadership.⁷ However, the significant overlap between these lists demonstrates an emerging national consensus on which technologies are most critical to U.S. national competitiveness. Table 5 illustrates the initial slate of technologies that the Commission recommends including as part of a broader technology leadership strategy, as well as whether or not those technologies have been included on select, existing U.S. government lists of critical technologies. As an initial step, the Commission recommends that the White House designate these technologies as critical through an Executive Order and direct Departments and Agencies to prioritize and coordinate them accordingly.

U.S. Government Lists of Critical Technologies						
NSCAI-Proposed Critical Technology List	2018 National Defense Strategy	DoD List of Critical Emerging Technologies	Commerce ANPRM on Emerging Technologies	PCAST List of Industries of the Future	S.3832 - Endless Frontier Act	WH Nat Strategy for C&ET
Artificial Intelligence	✓	✓	✓	✓	✓	✓
Biotechnology	✓	✓	✓	✓	✓	✓
Quantum Computing		✓	✓	✓	✓	✓
Semiconductors and Advanced Hardware	✓	✓	✓		✓	✓
Autonomy and Robotics	✓	✓	✓		✓	✓
5G and Advanced Networking		✓		✓	✓	✓
Advanced Manufacturing			✓	✓	✓	✓
Energy Systems	✓	✓			✓	✓

Actions to Promote Technologies and Platforms Essential to U.S. Leadership and National Security.

After reaching consensus on the set of emerging technologies essential to overall U.S. technology leadership, the Executive Branch should assess each sector and identify specific platforms that meet the following criteria:

- Have potential applications of strategic and national security importance;
- Could have a significant impact on overall U.S. technical leadership and competitiveness, either alone or when combined with existing U.S. technical strengths; and
- Require government action to spur or protect its development.

Such platforms could require government support for several reasons. In some instances, a market failure may lead to underinvestment by the private sector in an area of strategic importance to national security. In other instances, seizing a market opportunity may only be possible if the federal government focuses the private sector, academia, and research organizations on a specific goal. The government must tailor its approach to the context


by increasing funding, implementing regulatory changes, or taking other steps aimed at promoting innovation and protecting advantages that fit the circumstances.

The Commission has already presented recommendations to support U.S. leadership in key technology platforms within several of the aforementioned strategic technologies. For example, Chapter 11 of this report recommends establishing a National AI Research Resource, which would create an essential platform to sustain and extend U.S. leadership in AI. Additionally, in Chapter 13 of this report, the Commission provided a series of recommendations for promoting U.S. leadership in microelectronics, including specific actions to incentivize the construction of a leading-edge merchant fabrication facility domestically.

The recommendations below build on the Commission's previous work and provide further actions the U.S. government could take to promote U.S. leadership in the key associated technologies and platforms that the Commission assesses to be of greatest strategic importance—specifically, biotechnology, quantum computing, 5G and advanced networking, autonomy and robotics, advanced and additive manufacturing, and energy systems.⁸

Biotechnology.

Biology is now programmable, and AI's ability to identify ways to optimize this programming will enable transformational biotechnology breakthroughs. AI was crucial in the rapid development of COVID-19 vaccines, allowing researchers to finalize the genetic sequence of a vaccine candidate only two days after the virus' full genetic sequence was first posted online.⁹ Computer vision techniques applied to medical imagery have also enabled more accurate and efficient diagnoses.¹⁰ And recently, an AI network made substantial progress over the last year toward solving one of biology's most daunting challenges: determining a protein's 3D shape from its amino-acid sequence.¹¹ Tools such as these will become




“Biology is now programmable, and AI’s ability to identify ways to optimize this programming will enable transformational biotechnology breakthroughs.”

even more powerful in combination with synthetic biology and gene editing. Together they will enhance human health by allowing deeper studies of the building blocks of life and enabling the quicker discovery and fabrication of more advanced drugs and materials. As AI fuels rapid new developments in the biological sciences and biotechnology becomes a greater driver of the overall world economy, the strategic consequences of ceding leadership in biotechnology will increase significantly—a fact that the COVID-19 pandemic illustrates in clear and stark terms. The government should:

Recommendation

Prioritize the development of an advanced domestic biotechnology R&D ecosystem. As part of a national bioeconomy strategy, the United States should support the development of biotechnology platforms that maximize researchers' ability to utilize AI to drive new biological breakthroughs and help transition advanced research into physical products at scale. This will necessitate support for both world-class biodata resources to fully harness the power of AI and biomanufacturing platforms to rapidly realize the benefits from analytical breakthroughs:

- **Biodata:** The United States should fund and prioritize efforts to build a world-class biobank containing a wide range of high-quality biological and genetic data sets securely accessible by researchers. GenBank, the leading U.S. genetic database, which is run by the National Institutes of Health, is currently underfunded, underutilized, and poorly curated. The goal should be to create a genetic database that is well-curated and easy for researchers to access and use; contains a significant number and broad range of whole human, animal, and plant genomes; and aggregates open and proprietary data sets across government and the private sector. It should also contain de-identified metadata about corresponding phenotypes whenever possible and include strong privacy protections for human genetic data. This will require significant levels of funding; China National GeneBank, the equivalent facility in China that is operated by the BGI Group (formerly the Beijing Genomics Institute), required approximately \$117 million in initial funding.¹² Establishing such an entity in the United States would enhance and democratize biotechnology innovations by pooling existing data resources and facilitating new levels of AI-enabled analysis of genetic data while also reducing U.S. researchers' reliance on BGI or other Chinese entities for access to large-scale genomic databases for research.
- **Biomanufacturing:** The United States should support efforts to diversify and expand the biotechnology industry beyond its current vertically integrated models and encourage the development of multiple standardized, merchant biofabrication facilities. Doing so is necessary to ensure U.S. biomanufacturing capabilities keep pace with AI's transformative impact on the bioeconomy. Expanding access to advanced biofabrication tools among startups and laboratories would allow firms to rapidly design new molecules and materials via the cloud and place immediate orders for fabrication. Such efforts should include R&D funding and incentives to support advanced biotech manufacturing initiatives through entities such as the Biomedical Advanced Research and Development Authority (BARDA),¹³ with appropriate stewardship, and the expansion of existing relevant programs such as BioMADE.¹⁴ Congress should prioritize such initiatives in future health-related spending bills. Given that up to 60% of the physical inputs to the global economy could be produced via synthetic biology, there is a clear and pressing need for the United States to retain leadership in biomanufacturing moving forward.¹⁵



“... quantum computers have the potential to outperform their classical counterparts on certain classes of problems related to machine learning and optimization, the simulation of physical systems, and the collection and transfer of sensitive information.”

Quantum Computing.

As the pace of innovation predicted by Moore's Law becomes increasingly difficult for semiconductor manufacturers to maintain due to the physical limits of microchip design, leadership in next-generation computer hardware will be essential to preserving long-term U.S. advantages in strategic technologies like AI.¹⁶ Although classical computers will likely remain the most economical way of performing day-to-day computational tasks in the near future, quantum computers have the potential to outperform their classical counterparts on certain classes of problems related to machine learning (ML) and optimization, the simulation of physical systems, and the collection and transfer of sensitive information. For example, quantum computers may be able to efficiently optimize military logistics or discover new materials for weapon systems.¹⁷ Each of these applications creates novel national security threats and opportunities at the intersection of AI and quantum computing. The government should:

Transition from basic research to national security applications of quantum computing and incentivize domestic fabrication. The United States is a global leader in research of quantum computers, but it risks losing its edge in applications to national security. Recognizing that advances in quantum computing may drive advances in AI, the United States must establish trusted sources of materials and components for quantum computers, invest in the development of hybrid quantum-classical algorithms, and focus on fielding of national security applications. Offering access to both classical and quantum computers through

Recommendation

the National AI Research Resource will facilitate the development of hybrid quantum-classical algorithms that leverage noisy intermediate-scale quantum computers. Publicly announcing specific government use cases of quantum computers will signal that a market exists for national security applications and encourage further investment by the private sector.

5G and Advanced Networking.

5G networks will form the connective tissue between AI platforms, which means maintaining access to trusted and robust 5G networks is a critical component of overall leadership in AI. Huawei is pursuing global dominance in 5G, and there is no single supplier that can compete with it in terms of both price and quality. Due to the urgency of the issue, the United States should pursue several complementary approaches concurrently to ramp up deployment of 5G domestically and provide a credible alternative to Huawei. As a starting point, any comprehensive effort should include support for dynamic spectrum sharing.¹⁸ The government should:

“Expanding spectrum sharing efforts is critical to ensuring that DoD maintains access to spectrum essential for operational effectiveness while broadening commercial access to spectrum for 5G networks.”

Recommendation

Bolster and accelerate U.S. 5G network deployment through mid-band spectrum sharing. Expanding spectrum-sharing efforts is critical to ensuring that the Department of Defense (DoD) maintains access to spectrum essential for operational effectiveness while broadening commercial access to spectrum for 5G networks. A multi-agency effort is needed to expand sharing arrangements and licenses and permit additional portions of the mid-band to be simultaneously utilized by DoD and commercial carriers. Through this portfolio approach, the United States stands the best chance of accelerating its 5G deployment at a pace that can support the widespread adoption of AI.

Autonomy and Robotics.

Autonomous systems are already unlocking value across global markets. In the private sector, they enable products ranging from expert advisory systems and self-driving vehicles to manufacturing. In the realm of national security, autonomous systems generate opportunities to reduce the number of warfighters in harm's way, increase the pace and quality of decisions, and create entirely new military capabilities.¹⁹ The ability to design and produce the hardware and software for advanced robotics is an essential part of autonomous systems. The government should:

Incentivize the development of world-class software platforms for robotic and autonomous systems. The future of autonomy and robotics will manifest in almost unlimited shapes and sizes as firms develop and tailor robots for different use cases and environments. The U.S. trails nations such as China, Japan, and South Korea in the deployment of robots and robotic hardware and must work to improve its capabilities in such areas as materials design and energy storage for robots.²⁰ However, U.S. expertise in software development lends itself to creating a world-class digital platform for many classes of robotic hardware. The software powering robotic systems will be built upon several core capabilities rooted in AI: It will need to be able to sense its environment, reason, and operate in the world around it.²¹ In creating cutting-edge software for these types of capabilities, there is an opportunity for U.S. firms to win the market for the software platforms that power the next wave of industrialization.²² To promote U.S. leadership in the development of software for autonomous systems, the U.S. government should fuel industry's ongoing efforts by supplementing the basic R&D, standard-setting, and data-sharing programs led by National Institute of Standards and Technology (NIST)'s Intelligent Systems Division.²³ It should also incentivize early adoption of automation and create markets for autonomous systems in areas already ripe for them, such as mail sorting, that will yield data and experience relevant for achieving scale and addressing adjacent markets.²⁴ Combined, a multi-pronged approach along these lines would position industry to compete more effectively in the market for autonomous system software, a strategically important area aligned with existing U.S. technical strengths.

Recommendation

Advanced and Additive Manufacturing.

The capacity to produce high-tech goods domestically is critical to national security, both to maintain access to finished goods and as a driver of innovation. In terms of access, the United States must strive for self-reliance in industries that are critical to national security or that would take too long to regenerate in the event of protracted conflict.²⁵ Innovation also benefits from a tight feedback loop between technological design and production, which allows for more rapid iteration.²⁶ This link is particularly important in the defense sector, where feedback from the manufacturing process back into the R&D cycle helps bring technology from lab to military operations. Longer-term disruptions to the manufacturing industry through new techniques such as additive manufacturing also



“The capacity to produce high-tech goods domestically is critical to national security, both to maintain access to finished goods and as a driver of innovation.”

pose threats and opportunities for national security. For example, additive manufacturing may enable a step-change in domestic manufacturing capabilities, but it also creates new threats by potentially democratizing the production of firearms and other goods with military applications.²⁷ The government should:

Recommendation

Accelerate additive manufacturing production of legacy parts across the DoD. Additive manufacturing and 3D printing have the potential to transform manufacturing. They are capable of rapid, high-quality, and complex production, and they are flexible enough that 3D printers may be able to be located near the point of need for just-in-time production.²⁸ Although current additive manufacturing techniques struggle to replicate the quality of advanced traditional manufacturing techniques, AI has already shown the ability to enable significant improvements in their accuracy.²⁹ The Federal Government should proactively support initiatives that advance the development of additive manufacturing techniques and also provide practical benefits by easing the production of legacy items.³⁰ The DoD should announce a goal of identifying all legacy parts in active weapon systems that are capable of being produced via additive manufacturing and 3D printing and doing so by 2025.


Energy Systems.

Cheap and reliable access to energy is critical to U.S. national security, whether it be to ensure military readiness, facilitate the response to a domestic crisis, or keep the economy functioning smoothly. As an input to nearly every sector, the price of energy directly impacts economic output and is a key determinant of U.S. national competitiveness. Furthermore, dependence on foreign countries for energy resources and technologies would put the United States in a position of vulnerability, especially if those resources or technologies are controlled by strategic competitors. Although the United States is at the forefront of the exploration, extraction, and processing of oil and gas and possesses significant domestic

reserves, China is far and away the leading producer of renewable energy and is investing heavily in advanced energy storage technologies, such as batteries and their constituent materials.³¹ To remain competitive, in these critical sectors U.S. industry will need to achieve aggressive cost targets in terms of kilowatts/hour and energy density. This is especially true in markets with the most substantial growth potential, such as long-duration stationary storage devices and battery packs for electric vehicles.³² The government should:

Develop and domestically manufacture energy storage technologies to meet U.S. market demand by 2030. Developing new technologies to more effectively store electrical energy so it is readily available whenever and wherever needed would drive advances in electricity transmission and distribution. It would also offer advantages to the United States both economically and strategically. To accelerate breakthroughs in energy storage,³³ the Department of Energy has set the ambitious goal of developing and domestically manufacturing storage technologies capable of meeting the entirety of U.S. market demand by 2030.³⁴ Congress should fully fund the federal R&D and establish incentives for commercialization needed to achieve the Department of Energy's Energy Storage Grand Challenge roadmap by 2030.³⁵

Recommendation



“Cheap and reliable access to energy is critical to U.S. national security, whether it be to ensure military readiness, facilitate the response to a domestic crisis, or keep the economy functioning smoothly.”

Chapter 16 - Endnotes

¹ Recognizing this connection, Congress included AI and “associated technologies” as they relate to national security within the scope of the Commission’s mandate.

² The Commission’s first Interim Report identified biotechnology, quantum computing, and 5G as key emerging technologies associated with AI. See *Interim Report*, NSCAI at 50 (Nov. 2019), <https://www.nscai.gov/previous-reports/>.

³ Made in China 2025 includes the following sectors: new-generation information technology, high-grade machine tooling and robotics, aviation and aerospace equipment, marine engineering equipment and high-tech ships, advanced rail transportation equipment, new energy automobiles, electric power equipment, agriculture equipment, new materials, and biomedicine and high-tech medical devices. See Alice Tse & Julianna Wu, *Why ‘Made in China 2025’ Triggered the Wrath of President Trump*, South China Morning Post (Sept. 11, 2018), <https://multimedia.scmp.com/news/china/article/made-in-china-2025/index.html>.

⁴ According to McKinsey, “Platforms are the back-end technology capabilities, whether provided by individual systems or by assemblies of multiple systems, that power products.” Ross Frazier, et al., *Products and Platforms: Is Your Technology Operating Model Ready?*, McKinsey Digital (Feb. 28, 2020), <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/products-and-platforms-is-your-technology-operating-model-ready>.

⁵ The United States ranks 18th among 37 Organisation for Economic Co-operation and Development (OECD) countries in fixed and mobile broadband subscriptions per 100 inhabitants and eighth in average broadband speed. *Broadband Portal*, OECD (July 2020), <https://www.oecd.org/sti/broadband/broadband-statistics/> (see “Penetration and data usage” table “1.2 Fixed and mobile broadband subscriptions per 100 inhabitants” [Dec. 2019] and “Speeds” table “5.2 Akamai average speed” [Q1 2017]).

⁶ *Smart City Index*, IMD, 8 (Oct. 2019), [https://www.imd.org/research-knowledge/reports/imd-smart-city-index-2019/#:~:text=The%20Top%2010%20smartest%20cities,and%20Dusseldorf%20\(10th\)](https://www.imd.org/research-knowledge/reports/imd-smart-city-index-2019/#:~:text=The%20Top%2010%20smartest%20cities,and%20Dusseldorf%20(10th)).

⁷ For example, the White House published its National Strategy for Critical and Emerging Technologies in October 2020, which included a list of critical and emerging technologies as identified by Departments and Agencies. The document does not provide detail on how each technology is essential to national competitiveness and lacks a specific plan for promoting and protecting U.S. advantages in each. See *National Strategy for Critical and Emerging Technologies*, The White House at A-1 (Oct. 2020), https://www.whitehouse.gov/wp-content/uploads/2020/10/National-Strategy-for-CET.pdf?utm_source=morning_brew. In their 2018 report, Michael Brown and Pavneet Singh argue that the lack of a unified list of critical technologies harms the ability of the United States to protect against technology transfer. See Michael Brown & Pavneet Singh, *China’s Technology Transfer Strategy*, Defense Innovation Unit Experimental at 37 (Jan. 2018), [https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_\(1\).pdf](https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_(1).pdf).

⁸ This includes all technologies other than AI on NSCAI’s proposed critical technologies list with the exception of semiconductors, which are addressed separately in Chapter 13 of this report. Elements of AI-enabled biotechnology are also separately addressed in Chapter 1 of this report. Additional recommendations to promote U.S. leadership in biotechnology, quantum computing, and 5G can be found in the Chapter 16 Blueprint for Action.

⁹ Hannah Mayer, et al., *AI Puts Moderna Within Striking Distance of Beating COVID-19*, Harvard Business School (Nov. 24, 2020), <https://digital.hbs.edu/artificial-intelligence-machine-learning/ai-puts-moderna-within-striking-distance-of-beating-covid-19/>; Noah Weiland, et al., *Modern Vaccine Is Highly Protective Against Covid-19, the F.D.A. Finds*, New York Times (Dec. 18, 2020), <https://www.nytimes.com/2020/12/15/health/covid-moderna-vaccine.html>.

¹⁰ Junfeng Gao, et al., *Computer Vision in Healthcare Applications*, Journal of Healthcare Engineering (March 4, 2018), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5857319/>.

¹¹ Ewen Callaway, *‘It will Change Everything’: DeepMind’s AI Makes Gigantic Leap in Solving Protein Structures*, Nature (Nov. 30, 2020), <https://www.nature.com/articles/d41586-020-03348-4>.

¹² Zhuang Pinghui, *China Opens First National Gene Bank, Aiming to House Hundreds of Millions of Samples*, South China Morning Post (Sept. 22, 2016), <https://www.scmp.com/news/china/article/2021623/chinas-noahs-ark-first-national-gene-bank-opens-shenzhen>.

- ¹³ *Biomedical Advanced Research and Development Authority (BARDA)*, U.S. Department of Health and Human Services (Dec. 4, 2019), <https://www.hhs.gov/about/agencies/orgchart/aspr/barda/index.html>.
- ¹⁴ BioMADE is a public-private partnership with DoD, operated under Manufacturing USA, that is focused on building a sustainable bioindustrial manufacturing ecosystem and enhancing U.S. bioindustrial competitiveness. See *BioMADE (Bioindustrial Manufacturing and Design Ecosystem)*, Manufacturing USA (last accessed Jan. 9, 2021), <https://www.manufacturingusa.com/institutes/biomade>.
- ¹⁵ Michael Chui, et al., *The Bio Revolution*, McKinsey Global Institute at 43 (May 13, 2020), <https://www.mckinsey.com/industries/pharmaceuticals-and-medical-products/our-insights/the-bio-revolution-innovations-transforming-economies-societies-and-our-lives>.
- ¹⁶ Steve Blank, *What the GlobalFoundries' Retreat Really Means*, IEEE Spectrum (Sept. 10, 2018), <https://spectrum.ieee.org/nanoclast/semiconductors/devices/what-globalfoundries-retreat-really-means>.
- ¹⁷ Pontus Vikstål, et al., *Applying the Quantum Approximate Optimization Algorithm to the Tail-Assignment Problem*, Physical Review Applied, vol. 14, iss. 3 (Sept. 3, 2020), <https://doi.org/10.1103/PhysRevApplied.14.034009>; He Ma, et al., *Quantum Simulations of Materials on Near-term Quantum Computers*, npj Computational Materials (July 2, 2020), <https://doi.org/10.1038/s41524-020-00353-z>.
- ¹⁸ *The 5G Ecosystem: Risks & Opportunities for DoD*, DoD Defense Innovation Board (April 2019), https://media.defense.gov/2019/Apr/03/2002109302/-1/-1/0/DIB_5G_study_04.03.19.pdf.
- ¹⁹ *Summer Study on Autonomy*, DoD Defense Science Board (June 2016), <https://dsb.cto.mil/reports/2010s/DSBSS15.pdf>.
- ²⁰ Maximiliano Dvorkin & Asha Bharadwaj, *Which Countries and Industries Use the Most Robots*, Federal Reserve Bank of St. Louis (Nov. 7, 2019), <https://www.stlouisfed.org/on-the-economy/2019/november/robots-affecting-local-labor-markets>.
- ²¹ Advanced materials (such as biological components), brain-computer interfaces, and small and efficient power supplies are additional areas of potential innovation that connect robotics to AI and other associated technologies described in this chapter.
- ²² For example, core capabilities might include gripping physical objects, which robotics maker ABB and several firms in the U.S. and Europe are currently pursuing. See Jonathan Vanian, *Industrial Robotics Giant Teams Up with a Rising A.I. Startup*, Fortune (Feb. 25, 2020), <https://fortune.com/2020/02/25/industrial-robotics-ai-covariant/>.
- ²³ *Intelligent Systems Division*, NIST (last accessed Jan. 6, 2021), <https://www.nist.gov/el/intelligent-systems-division-73500>.
- ²⁴ One specific area to expand demand for autonomous systems could be drastically scaling the U.S. Postal Service's Autonomous Mobile Robot pilot program from 25 sorting facilities to all sorting facilities by 2025. *Autonomous Mobile Robots and the Postal Service*, USPS Office of Inspector General (April 9, 2018), <https://www.uspsoig.gov/sites/default/files/document-library-files/2019/RARC-WP-18-006.pdf>.
- ²⁵ *Critical Technology Accessibility*, National Academies Press (2006), <https://www.nap.edu/read/11658/chapter/1>; see also *Assessing and Strengthening the Manufacturing and Defense Industrial Base and Supply Chain Resiliency of the United States*, Interagency Taskforce in Fulfillment of Executive Order 13806 at 46 (Sept. 2018), <https://media.defense.gov/2018/oct/05/2002048904/-1/-1/1/assessing-and-strengthening-the-manufacturing-and%20defense-industrial-base-and-supply-chain-resiliency.pdf> (identifying 10 risk archetypes threatening America's manufacturing and defense industrial base).
- ²⁶ *Strategy for American Leadership in Advanced Manufacturing*, National Science & Technology Council (Oct. 2018), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2018/10/Advanced-Manufacturing-Strategic-Plan-2018.pdf>; Gregory Tassej, *Rationales and Mechanisms for Revitalizing US Manufacturing R&D Strategies*, NIST (Jan. 29, 2010), https://www.nist.gov/system/files/documents/2017/05/09/manufacturing_strategy_paper_0.pdf.

Chapter 16 - Endnotes

²⁷ *3D Opportunity for Adversaries*, Deloitte (Aug. 22, 2017), <https://www2.deloitte.com/us/en/insights/focus/3d-opportunity/national-security-implications-of-additive-manufacturing.html>.

²⁸ *Audit of DoD's Use of Additive Manufacturing for Sustainment Parts*, DoD Inspector General (Oct. 17, 2019), <https://media.defense.gov/2019/Oct/21/2002197659/-1/-1/1/DODIG-2020-003.pdf>.

²⁹ Mark Anderson, *3D Print Jobs Are More Accurate with Machine Learning*, IEEE Spectrum (Feb. 19, 2020), <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/3d-print-jobs-news-accurate-machine-learning>.

³⁰ For instance, in August 2020, the DoD printed the first metal part for a B-52 jet engine. Kyle Mizokami, *The Old School Engine That Powers the B-52 Gets a 3D-Printed Upgrade*, Popular Mechanics (Aug. 10, 2020), <https://www.popularmechanics.com/military/aviation/a33535790/air-force-3d-print-metal-part-turbofan-engine/>.

³¹ Robert Rapier, *Ten Countries That Dominate World Fossil Fuel Production*, Forbes (July 14, 2019), <https://www.forbes.com/sites/rrapier/2019/07/14/ten-countries-that-dominate-fossil-fuel-production; Country-Rankings>, International Renewable Energy Agency (last accessed Jan. 6, 2021), <https://www.irena.org/Statistics/View-Data-by-Topic/Capacity-and-Generation/Country-Rankings>.

³² *Energy Storage*, U.S. Department of Energy (last accessed Jan. 6, 2021), <https://www.energy.gov/oe/energy-storage>.

³³ The field of energy storage includes a broad technology base such as batteries (both conventional and advanced), electrochemical capacitors, flywheels, power electronics, control systems, and software tools for storage optimization and sizing.

³⁴ *Energy Storage*, U.S. Department of Energy (last accessed Jan. 6, 2021), <https://www.energy.gov/oe/energy-storage>.

³⁵ *Energy Storage Grand Challenge: Roadmap*, U.S. Department of Energy (Dec. 2020), <https://www.energy.gov/sites/prod/files/2020/12/f81/Energy%20Storage%20Grand%20Challenge%20Roadmap.pdf>.

Blueprints for Action PART ONE



Chapter 1: Emerging Threats in the AI Era	273
Combating Malign Information Operations Enabled by AI	273
Preparing for AI-Enabled Cyber Conflict	278
Chapter 2: Foundations of Future Defense	291
Chapter 3: AI and Warfare	337
Chapter 5: AI and the Future of National Intelligence	349
Chapter 6: Technical Talent in Government	357
Chapter 7: Establishing Justified Confidence in AI Systems	379
Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security	395

The following Blueprints for Action cover Part I of NSCAI’s Final Report. Part I, “Defending America in the AI Era,” (Chapters 1-8) outlines what the United States must do to defend against the spectrum of AI-related threats from state and non-state actors, and recommends how the U.S. government can responsibly use AI technologies to protect the American people and our interests. These Blueprints for Action complement the Commission’s Final Report and mirror its organizational structure.

The Blueprints for Action serve as more detailed roadmaps for Executive and Legislative branch actions to retain America’s AI leadership position. They identify who should take a particular action—Congress, the White House, or an executive branch department or agency. The Commission provides estimated increases in funding or appropriations as part of its recommendations. All recommendations that include funding figures should be considered estimates for consideration by Congress and/or the Executive branch.

This report does not contain a separate Blueprint for Action for Chapter 4. This is because given the importance of the topic, the Commission chose to detail its arguments, recommendations, and the specific actions required to implement them directly in this chapter. Additionally, further detail on how the United States should adapt its TEVV policies to maintain confidence in AI systems can be found in Chapter 7 and its associated Blueprint for Action, and recommendations on relevant changes to DoD organizational structure can be found in Chapter 3.

Chapter 1: Emerging Threats in the AI Era

Blueprint for Action: Number One

Combating Malign Information Operations Enabled by AI.

The use of AI to produce, manipulate, and promote malign information marks a disruptive evolution in the use of information as a tool of statecraft, a weapon of war, and a threat to democracy.¹ The following recommendations represent a strategic, organizational, and operational framework that the U.S. government should adopt to adequately defend and counter malign information operations in the AI era, including by employing AI-enabled technologies.

Recommendation

Recommendation: A National Strategy for the Global Information Domain

Expanding upon the principles of information statecraft outlined in the 2017 National Security Strategy,² the President should issue a new national strategy for the global information domain that more fulsomely addresses how AI and associated technologies are defining new fronts in this area. The strategy should:

- Acknowledge that the network-connected world is dissolving barriers between societies.
- Prioritize the global information domain as an arena for competition.
- Detail how adversarial state and non-state actors are attempting to define and control the global information domain in order to shape global opinion and achieve strategic advantage.
- Account for the critical role of AI-enabled malign information in achieving these goals.
- Designate malign information operations as a national security threat with its own set of priority actions to defend, counter, and compete against them.
- As necessary, update critical infrastructure designations and require relevant departments and agencies to update sector-specific plans to reflect emerging technologies.
- Establish organizational structures for U.S. national security agencies to defend, counter, and compete against the threat.

Action for the President:

- Issue a supplemental National Strategy for the Global Information Domain.

Action for Congress:

- Congress should direct the Executive Branch to transmit a National Strategy for the Global Information Domain that categorizes the global information domain as an arena of competition vital to the national security of the United States.

Organizational Framework

The proliferation of malign information has exposed an Achilles heel in the U.S. national security apparatus. Previous major reorganizations could not foresee contemporary digital technology and society's profound dependence upon it. They could not anticipate the use of ICT platforms and tools, bots, and AI-enabled technologies to spread false information. They do not account for the role that the commercial sector and civil society play in defending against malign information, and enabling its spread. Individual agencies such as the Department of Homeland Security (DHS), the Federal Bureau of Investigation (FBI), and the Office of the Director of National Intelligence (ODNI) have stretched their mandates to confront the threat. They rely on narrow sets of outdated tools, and are hampered by cultures shaped by the Cold War and counter-terrorism paradigms.

Recommendation: Create a Joint Interagency Task Force (JIATF) and Operations Center.

Recommendation

Action for the President:

- **Direct creation of a JIATF and operations center to lead and integrate government efforts to counter foreign-sourced malign information in real time.**
 - o The Presidential action should direct the Secretaries of State, Defense, Justice, and Homeland Security, the Attorney General and the Director of National Intelligence, to create a JIATF and stand-up an operations center to counter foreign-sourced malign information.
 - o The JIATF should integrate efforts of key offices, bureaus, and divisions within each of these agencies, as well as the broader intelligence community (IC) and law enforcement establishment.
 - o The JIATF should have the responsibility to survey the landscape of relevant public and private actors, coordinate among them, and act in real time to counter foreign information campaigns.
 - o The JIATF should draw on existing authorities to create an operations center with modern, AI-enabled digital tools and expert staff to expose, attribute, and respond effectively.
 - o The Presidential action should also direct these officials, as part of the JIATF, to create a mechanism to share and exchange critical information with key companies in the private sector that run internet and social media platforms where malign information proliferates.

Action for the Secretaries of State, Defense, Justice, Homeland Security, and the Director of National Intelligence:

- **Establish the JIATF and Operations Center.**

- o These agency heads should direct immediate development of a plan to create the JIATF and operations center with a focus on identifying those offices, bureaus, and divisions within their agencies and the broader IC and law enforcement establishment that are essential to the mission of countering foreign-sourced malign information.
 - As part of this effort, the JIATF should leverage the authority provided by Congress in the FY2020 NDAA to stand-up a Foreign Malign Influence Response Center within ODNI.³
 - Components that will be critical to the JIATF include, among others, the Central Intelligence Agency's Open Source Enterprise and the National Counterintelligence and Security Center.⁴ Leadership will need to ensure involvement of relevant components from the FBI, the National Security Agency, across the Department of Defense, and the Global Engagement Center (GEC) at the Department of State.
- The JIATF would lead and integrate existing and new national strategic efforts against foreign malign information operations by providing analysis, sharing information with government and commercial partners, and driving whole-of-government *action*, subject to Presidential direction, to advance U.S. information objectives.
- The Commission proposes that the operations center component of the JIATF be modeled on the National Counterterrorism Center (NCTC), as a proven model for providing real-time situational awareness of and response to evolving national security threats.
- To exchange information and coordinate with internet and social media platforms on malign information threats, the Commission proposes creation of an associated industry consortium that includes an information sharing and analysis center (ISAC). The consortium, supplemented by the ISAC, would allow the JIATF to exchange information *with industry*, monitor malign information across ICT platforms, and improve U.S. government response to malign information threats. In developing the ICT consortium and ISAC, JIATF should look to the Global Internet Forum to Counter Terrorism as a model.⁵

Action for the Director of National Intelligence:

- **Appoint a Malign Information Threat Executive (MITE) to lead the JIATF.**

- o In July 2019, ODNI created the Election Threat Executive position responsible for coordinating across the IC on issues related to election security.⁶ The threat of foreign malign information operations demands that this position be elevated, renamed, and expanded beyond the subject of elections.
- o The MITE role should also serve as a liaison function between the White House/ National Security Council and the JIATF to ensure alignment and responsiveness to the national security strategy.

Action for Congress:

- **Appropriate \$30 million per year to support the operations of the JIATF.**

Operational Framework

Efforts by the U.S. Government and private sector to counter terrorist propaganda offer a potential roadmap for how the United States can go on the offensive to counter and compete against malign information. The creation of the Global Coalition to Defeat the Islamic State of Iraq and Syria (ISIS) has shown how a burden-sharing model can be deployed to successfully counter and defeat a shared threat.⁷ The United States and its allies will only succeed if they can develop and deploy personnel as well as an advanced set of tools to assist in their effort to counter and compete against malign information operations. Efforts need to be made to encourage innovation as well as harness commercially available technologies to go on the offensive.

Recommendation: The Department of State should lead a global effort to counter disinformation.

Recommendation

Action for the President:

- **Designate the Under Secretary of Public Diplomacy and Public Affairs at the Department of State to lead the international fight against malign information operations.**

Action for the Department of State:

- **Build an International Task Force to Counter and Compete Against Disinformation.** Modeled after the Global Coalition to Defeat ISIS, the Department of State should build a similar task force to counter malign information. The International Task Force to Counter and Compete Against Disinformation should be led by the Department of State's Under Secretary for Public Diplomacy and Public Affairs, with the GEC coordinating its daily activities.⁸ The task force will be in charge of directing, leading, synchronizing, integrating, and coordinating efforts by allies to recognize, understand, expose, and counter foreign state and non-state propaganda and malign information efforts. The GEC should leverage the work of the Technology Engagement Team (TET) to share and test technologies to detect and disrupt the creation, manipulation, and dissemination of malign information from state and non-state actors. *See the Chapter 15 Blueprint for Action for more detail on creating a task force as part of the Emerging Technology Coalition proposed by the Commission.*

Recommendation: The Defense Advanced Research Projects Agency (DARPA) should coordinate multiple research programs to detect, attribute, and disrupt AI-enabled malign information campaigns and to authenticate the provenance of digital media.

Recommendation

The government should sponsor research to develop technologies to detect, attribute, and disrupt malign influence operations, including influence campaigns, psychological operations on social media platforms, and manipulated and synthetic media. In parallel,

the government should develop alternative technologies to authenticate the provenance of digital media and head off the risk that other approaches will not be successful. These efforts should be led by DARPA.

Action for Congress:

- **Appropriate \$60 million to \$80 million in additional funding for DARPA to sponsor multiple research projects to develop technologies to detect, attribute, and disrupt malign influence operations that rely on AI-generated content, and to develop alternative technologies to authenticate the provenance of digital media.**⁹ DARPA has existing authority to fund such research with the scope outlined in this recommendation, but will require dedicated appropriations to carry out the effort and a security review of the best innovation vehicles to sponsor the research.

Action for DARPA:

- **Sponsor further research as described above using innovation vehicles, such as challenge competitions, or any other deemed necessary by DARPA to develop and transition these technologies to accountable agencies and departments for maximum employment.**

Recommendation

Recommendation: Create a task force to study the use of AI and complementary technologies, including the development and deployment of standards and technologies, for certifying content authenticity and provenance.

In response to the challenges of misinformation, efforts are underway to develop standards and pipelines aimed at certifying the authenticity and provenance of audiovisual content.¹⁰ These efforts make use of technologies, including encryption and fragile watermarking, to secure and track the expected transformations of content via production and transmission pipelines. These efforts offer the opportunity to mitigate malign information campaigns that seek to corrupt or spoof highly trusted sources of information across our digital ecosystem. This technology area is ripe for public-private partnership, as several private organizations are already forming to fight disinformation.¹¹

Actions for the Office of Science and Technology Policy (OSTP):

- **Establish a task force to study the use of AI and complementary technologies for certifying content authenticity and provenance.**
 - o OSTP should establish an interagency task force to assess the use of AI and complementary technologies to certify content authenticity and provenance, to include an evaluation of technical standards and production and transmission pipelines.
 - o The task force should make recommendations on methods to improve content certification, which may include public-private initiatives, legislation, and changes to federal policy. In addition, the task force should assess options for federal regulation of content certification by non-governmental organizations.

Recommendation: Executive Branch departments and agencies should utilize Other Transaction Authorities (OTAs), creative investing, and the Small Business Innovation Research (SBIR) program to deploy capital to companies that offer technical solutions that will assist the United States Government in identifying, countering, and defending against malign information operations.

Recommendation

The U.S. Government has an array of mechanisms that are not currently leveraged to deploy capital to companies that create strategic technology to unleash AI, machine learning (ML), and associated technologies in this counter-information operations fight.¹²

Action for all U.S. departments and agencies:

- **Explore the use of the SBIR program and OTAs to acquire technology solutions that will assist the United States Government in identifying, countering, and defending against malign information operations.**

The United States must prepare for both the present and future threat of increasingly automated and AI-enabled cyber conflict. The expanding threats of mutating malware and AI-powered tools are combining with traditional cyber threats to automate, optimize, and ultimately transform the precision, speed, stealth, scale, and effectiveness of cyber-attack and espionage campaigns.¹³ To defend the U.S. from current and future cyber threats, we must move to develop AI-enabled cyber defenses and to mitigate proliferating cyber vulnerabilities.

Chapter 1: Emerging Threats in the AI Era

Blueprint for Action: Number Two

Preparing for AI-Enabled Cyber Conflict.

Section 1: Developing AI-enabled defenses against cyber attacks.

Recommendation: Develop and deploy machine-speed threat detection and mitigation.

Recommendation

Detecting and reacting to unknown threats on a network is difficult, but not impossible, for self-learning AI systems that have been trained to differentiate between normal and anomalous network behavior.¹⁴ To address deficiencies highlighted by the SolarWinds attack, autonomous defenses are needed to defend the U.S. Government's systems.

Actions for the Department of Homeland Security and Department of Defense:

- **Expand machine speed threat information sharing, behavior-based anomaly detection, and cyber threat mitigation to all government networks containing sensitive information and critical functions.**
 - o DHS must improve the National Cybersecurity Protection System (NCPS), while DoD must also accelerate its efforts to harness AI-enabled cyber defenses and sensors. At a minimum, the objective of these new defenses should be to flag or potentially block never-before-seen connections and communications missed by currently deployed intrusion detection and prevention technologies such as EINSTEIN.¹⁵ To fully take advantage of new capabilities, these defenses should also aim to accelerate recovery from cyber attack by automatically generating courses of action for federal agencies to assure secure continuity of operations. These defenses should assist recognition of insider threats as well as externally launched attacks, and use machine speed information sharing to prepare other public and private networks to defend themselves against detected threats.
 - o DoD and DHS must also assess and mitigate security risks posed by introducing and enhancing threat detection systems. These systems will require precautions against their elevated system access being used to deliver malware or abused by other cyber threats. AI-enabled system components designed to mitigate new and unknown threats likewise will need defenses against adversarial techniques.
 - o To minimize cost overruns in altering a multibillion-dollar project, DHS should reprogram \$10 million to investigate the best means to accelerate and set up AI-enabled threat detection systems. This study would be tasked to look for synergies with existing intrusion detection software and infrastructure, seek to address any remaining key deficiencies found by GAO in the National Cybersecurity Protection System, and to develop a final budget proposal for Congress.¹⁶ This study likewise should aim to address how previous intrusion detection systems failed to detect the SolarWinds cyber attack.

Recommendation

Recommendation: Execute large, instrumented, and realistic tests to gather data and train AI-enabled cyber defenses.

AI-enabled cyber defenses require training to recognize potential threats, and sensors to detect them. By experimenting with larger networks in realistic conditions, the United States can train more robust AI-enabled cyber defense capabilities.

Action for Congress:

- **Fund the Defense Advanced Research Projects Agency (DARPA) to sponsor additional secure, instrumented, and realistic research on AI-enabled cyber defenses.**
 - o DARPA funding should be increased by \$20 million, to be divided between a security review, and other programmatic costs for the additional research. DARPA should be left free to determine the structure of further research, with an innovation vehicle such as a challenge competition or any other that DARPA deems necessary.
- **Expand the National Institute of Standards and Technology AI testbed program.**

- o For FY2021, NIST requested a \$25 million increase, for measurement tools and testbeds to accelerate the development and adoption of interoperable, secure, and reliable AI technologies.¹⁷ Since then, NIST has been authorized for \$64 million in additional AI R&D responsibilities including AI testbeds. To ensure NIST can meet its new responsibilities in addition to its prior ones, Congress should meet NIST's authorized funding increase for AI R&D.

Actions for DARPA:

- **Structure and standardize an innovation vehicle, such as a challenge competition, or any other DARPA deems necessary, to increase insight about options for new AI-enabled cyber defenses.**
 - o DARPA should aim to encourage the prototyping of new means of AI-enabled cyber defense and test the efficacy of these defenses against intelligent opponents and AI-enabled cyber threats. DARPA should structure new research to broaden insight on the importance of real-life factors such as cyber-attack externalities, differences in risk tolerance between threat actors, and differences in network infrastructure between defenders.¹⁸
- **Bring broader fields of expertise to bear for cyber defense research.**
 - o Cyber expertise is not the only expertise relevant to cybersecurity and the efficacy of cyber operations.¹⁹ The new research should involve experts from other fields such as economics, game theory, and behavioral psychology to improve scoring metrics, improve the human components of cyber strategy, and propagate insight further within government. With these improved metrics and insights, future investments can be more directly aligned with mission assurance.
- **Conduct a security review to determine the rules and bounds of new cyber research initiatives.**
 - o DARPA must conduct a thorough security review about the second-order effects of sponsoring research with public-facing results and without strong information security measures, to mitigate against potential adversaries acquiring information that can be weaponized against us. International competition in this area is getting so intense that the organization must consider using a vetted closed-challenge competition or initiative as opposed to an open-challenge competition format.

Actions for NIST:

- **Expand the NIST AI testbed program to generate data for AI-enabled cyber defenses in differing IT infrastructure environments.**
 - o Larger-scale testing is necessary to generate the data required for AI-enabled cyber defenses. By scaling testbeds within NIST, there will be the opportunity to generate this data, and to evaluate the performance of varying network architectures at strengthening network security.
 - o Training data often reflects a broad sampling of common scenarios and does not itself necessarily convey the costs of different types of compromises without further labeling.²⁰ *NIST should create optimized data sets for training cyber defenses to minimize expected costs of network disruption, compromise, and damage* rather than merely trying to identify cyber threats and vulnerabilities with high accuracy. To develop these data sets, NIST will need to hire or contract multidisciplinary talent to develop better metrics.

Recommendation

Recommendation: Ensure the robustness of AI-cyber defenses.

To make AI-based cyber defenses stronger, their supporting supply chains and data must be defended, while the algorithms themselves must be protected from malware, trained against adversarial techniques, and red teamed to the point of failure. *This approach can be found in the Chapter 7 Blueprint for Action.*

Section 2: Ensuring resilience against AI-enabled cyber attacks.

Many of the defenses required to protect against AI-enabled cyber threats are also required to defend against less advanced cyber threats. To provide this protection, the Commission endorses specific Cyberspace Solarium Commission recommendations, which are instrumental in enhancing U.S. defenses against AI-enabled cyber threats.²¹

Recommendation: Improve incentives for information and cyber security.

Recommendation

AI cannot defend inherently indefensible digital infrastructure against escalating offensive AI-enabled cyber capabilities. Even if vulnerabilities are known and easily patchable, that is no guarantee that they will be closed without a further impetus to action. Similarly, while new instrumented digital infrastructure is required to accelerate AI-enabled cyber defenses, those that build it must be careful to ensure new vulnerabilities don't outweigh the benefits of these defenses. In both cases, incentives must be realigned in the public and private sector to assure gaps are closed and new infrastructure is secure.

Action for Congress:

- **Establish liability for final goods assemblers for damage stemming from incidents that exploit known and unpatched vulnerabilities, incentivize reporting, and amend the Sarbanes-Oxley Act to include cybersecurity reporting requirements.**²²
 - o The Cyberspace Solarium Commission made recommendations to incentivize timely vulnerability patching. In addition to these recommendations, companies should be incentivized to improve their cybersecurity, and participate in new vulnerability disclosure programs via selectively reducing legal liability and product recalls for companies that can mitigate and patch controlled vulnerabilities within a limited, but rule-defined, time period. The overall structure of liability reform should aim to minimize perverse incentives to avoid liability by concealing failure. Grid, critical infrastructure, and medical device companies should be the primary targets for improvement.
 - o To harmonize and clarify cybersecurity oversight and reporting requirements for publicly traded companies, Congress should amend the Sarbanes-Oxley Act to explicitly account for cybersecurity.²³

Action for the Executive Branch

- **Incentivize information technology security through Federal Acquisition Regulations and Federal Information Security Management Act authorities.**²⁴
 - Zero-trust networking and robust code should become key priorities for government contracts related to information technology, and especially for contracts related to AI. Contractors should not be paid more for additional lines of code when adding them generates new vulnerabilities without additional functionality. Code should be subjected to AI-enabled vulnerability review.
- **Task CISA to develop an IT infrastructure “Cash for Clunkers” incentive plan, to submit to Congress for FY2022.**
 - This program would support the replacement of vulnerable outdated equipment with modern alternatives through targeted federal subsidies. CISA should coordinate the effort by setting the program’s strategy, prioritizing devices and critical digital infrastructure for replacement, and determining subsidy levels for the systems to be replaced. CISA must develop the plan so as to minimize perverse incentive to acquire vulnerable infrastructure before the plan is funded, and once the plan is developed, Congress must implement it as quickly as possible to reduce perverse incentives for companies to hold out on replacing vulnerable devices and infrastructure in the meantime.

Section 3: Disrupting adversary AI-enabled cyber-attacks and capabilities.

Recommendation: Develop additional, impactful non-kinetic options to respond to adversarial cyber and information operations.

Recommendation

Modern information operations have enormous overlap with cyber operations. As AI-enabled cyber capabilities spread in the presence of wide-open societal vulnerabilities, the United States needs to have additional tools to counter proliferating threat actors, and to establish deterrence in the cyber and information domains.

Action for Congress:

- **Expedite the establishment of the Bureau of Cyberspace Security and Emerging Technologies (CSET) within the U.S. Department of State.**
 - The CSET Bureau will be essential for strengthening norms in cyberspace, engaging other countries on information technology standards, assisting allied cyber defense, and improving international cyber law enforcement. *Recommendations to expedite the Bureau’s buildout and ensure that it has a clear mandate to coordinate strategy on the full range of emerging technology issues, in addition to critical cybersecurity needs, can be found in the Chapter 15 Blueprint for Action.*
- **Strengthen the U.S. Government’s ability to take down botnets by enacting Section 4 of the International Cybercrime Prevention Act.**²⁵

- o Botnets are already a present threat, and may become more powerful with advances in AI, not just directly spreading malware, but harvesting both computational power and data to put toward further offensive training in ways that were not previously possible. To enable the U.S. Government to better work with private industry and international partners, Congress, in consultation with the Department of Justice, should enact Section 4 of the International Cybercrime Prevention Act.²⁶ This legislation would provide broader authority to disrupt all types of illegal botnets, not just those used in fraud.²⁷

Actions for Cyber Command, the Department of Homeland Security, the Federal Bureau of Investigation, and the National Security Agency:

- **Expand current cyber threat inoculation initiatives.**

- o Machine speed information sharing is a key piece of enabling AI-cyber defenses. To contribute to the readiness of U.S. defense and critical infrastructure, efforts should be made to accelerate sharing of the most recent malicious code captured in the wild through appropriate interagency channels, including through a Joint Collaborative Environment. U.S. Cyber Command should ensure and accelerate coordination with DHS, the FBI, NSA, and stakeholders in the private sector in the release of threat information, particularly with owners and operators of systemically important critical infrastructure.²⁸

Section 4: Coordinating and Strategizing a Response.

Recommendation

Recommendation: Reform the U.S. Government's strategy, structure, organization, and authorities for handling AI-enabled cyber threats.

The U.S. must organize and align authorities to fully implement the cyber security mission and fully capitalize on machine speed information sharing defenses. Technology alone isn't enough: Cyber threat intelligence, joint planning, and response must be integrated into the same organization to keep pace with AI cyber threats.

Actions for the Executive Branch:

- **Issue an updated National Cyber Strategy with the following components.**

- o First, the strategy should build on the layered deterrence framework put forward by the Cyberspace Solarium Commission with a focus on making the framework more robust against the ways AI will transform cyber conflict.²⁹
 - To support the strategy, the Department of Defense, in partnership with the Department of State and the IC, should also develop a multitiered signaling strategy and promulgate a declaratory policy that addresses the use of AI in cyber operations.³⁰
- o Second, to inform the strategy, the Department of Homeland Security should run a study to develop regulatory recommendations for the most cost-effective means of defending digital devices and infrastructure. This study should investigate, but not be limited to:

- Standards requiring critical private and public sector networks to keep their data encrypted at rest and in transit
 - Multifactor authentication requirements for critical private and public sector networks
 - Air gapping requirements for select sensitive, but still unclassified, networks
 - Analog defenses for cyber physical infrastructure to prevent the most lethal failures regardless of how much network access cyber attackers gain, or how advanced their methods of attack become
 - Federated machine learning techniques that lower espionage and privacy risk via enabling data to be partitioned or remain decentralized
 - Specialized, narrow purpose computation hardware that can't be repurposed by malware for attacks
 - Ways to harness AI to lock down and constrain hardware toward its intended purpose on vulnerable networks that can't yet be patched or replaced
 - Ways to use cloud computing and virtual machines to reduce vulnerability of AI and cyber systems to advanced persistent threats
- **Accelerate the establishment of a Joint Cyber Planning and Operations Center, modeled after the National Counterterrorism Center.**³¹
 - This planning office under the Cybersecurity and Infrastructure Security Agency is necessary to coordinate cybersecurity planning and readiness across the federal government and between public and private sectors. To properly stand-up such a collaborative environment, the Executive branch must submit to Congress a list of authorities and data sharing issues that will require additional authorities or funding.
 - **Develop and implement an information and communications technology industrial base strategy.**³²
 - This strategy must increase support to supply chain risk management efforts, and provide better defense to the hardware supply chains, data, and algorithms that compose the "AI stack."

Action for Congress:

- **Establish a Bureau of Cyber Statistics to inform both cyber defense policy and AI-enabled cyber defenses.**³³
 - Large accurate data sets with relevant data are especially useful for training AI-enabled cyber defenses that minimize the costs of cyber attacks and false alarms, rather than just the number of attacks and false alarms. To that end, Congress should establish a Bureau of Cyber Statistics, within the Department of Commerce, or another department or agency, that would act as the government statistical agency that collects, processes, analyzes, and disseminates essential statistical data on cybersecurity, cyber incidents, and the cyber ecosystem to the American public, Congress, other federal agencies, state and local governments, and the private sector.³⁴

Recommendation

Recommendation: Coordinate with the Private Sector to Increase Resilience Against AI-Enabled Cyber Attacks.

Action for Congress:

- **Create or Designate Critical Technology Security Centers.**³⁵
 - o Congress should direct and appropriate funding for the Department of Homeland Security, in partnership with the Department of Commerce, Department of Energy, Office of the Director of National Intelligence, and Department of Defense, to competitively select, designate, and fund up to three Critical Technology Security Centers.
 - o These Centers would be designed to centralize efforts directed toward evaluating and testing the security of devices and technologies that underpin our networks and critical infrastructure.

- **Authorize, establish, and fund a joint collaborative environment for sharing and fusing threat information.**³⁶
 - o Sharing and fusing threat information is an instrumental step in improving the speed and capability of potential AI-enabled cyber defenses.³⁷ Congress must ensure that Executive branch agencies have necessary authorities to bring their data together in support of these efforts. Likewise, Congress must create incentives—including liability protection—to attract the private sector to participate in threat information sharing programs.
 - o To achieve these goals, the Commission endorses the Cyberspace Solarium Commission recommendation for Congress to establish a 'Joint Collaborative Environment,' a common, cloud-based environment in which the federal government's unclassified and classified cyber threat information, malware forensics, and network data from monitoring programs are made commonly available for query and analysis—to the greatest extent possible.³⁸

Blueprint for Action: Chapter 1 - Endnotes

¹ For the purposes of this section, “malign information” includes both disinformation—false information or intentionally misleading facts communicated with the intent to deceive—and misinformation—false information not necessarily meant to deceive. See Daniel Fried & Alina Polyakova, *Democratic Defense Against Disinformation*, Atlantic Council at n.1 (Feb. 2018); https://www.atlanticcouncil.org/wp-content/uploads/2018/03/Democratic_Defense_Against_Disinformation_FINAL.pdf. For a broader discussion, see Laura Rosenberger, *Making Cyberspace Safe for Democracy*, Foreign Affairs (May/June 2020), <http://www.foreignaffairs.com/articles/china/2020-04-13/making-cyberspace-safe-democracy>. For a study of how AI might be used to counter disinformation, see William Marcellino, et al., *Human-machine Detection of Online-based Malign Information*, RAND Europe (2020), https://www.rand.org/content/dam/rand/pubs/research_reports/RRA500/RRA519-1/RAND_RRA519-1.pdf.

² *National Security Strategy of the United States*, The White House at 34 (Dec. 18, 2017), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2017/12/NSS-Final-12-18-2017-0905-2.pdf>.

³ Pub. L. 116-92, The National Defense Authorization Act for FY2020, 133 Stat. 1198, 2129-30 (2019).

⁴ *National Counterintelligence and Security Center*, Office of the Director of National Intelligence (last accessed Feb. 8, 2020), <https://www.dni.gov/index.php/ncsc-home>.

⁵ *About*, Global Internet Forum to Counter Terrorism (last accessed Oct. 2, 2020), <https://www.gifct.org/about/>.

⁶ Press Release, Office of the Director of National Intelligence, *Director of National Intelligence Daniel R. Coats Establishes Intelligence Community Election Threats Executive*, (July 19, 2019), <http://www.dni.gov/index.php/newsroom/press-releases/item/2023-director-of-national-intelligence-daniel-r-coats-establishes-intelligence-community-election-threats-executive>.

⁷ Brett McGurk, *America Should Build an International Coalition Now*, The Atlantic (March 29, 2020), <https://www.theatlantic.com/ideas/archive/2020/03/america-should-build-international-coalition-now/608983/>.

⁸ Though this overall Blueprint for Action uses the term “malign information” to broaden beyond disinformation to include misinformation, it will likely be easier to organize a task force around countering disinformation.

⁹ Funding level should depend upon the number of programs DARPA has the capacity to execute in this area.

¹⁰ See, e.g., Paul England, et al., *AMP: Authentication of Media via Provenance*, arXiv (June 20, 2020), <https://arxiv.org/abs/2001.07886>.

¹¹ See *Creating the Standard for Digital Content Attribution*, Content Authenticity Initiative (last accessed Feb. 19, 2020), <https://contentauthenticity.org/>; *Overview*, Project Origin (last accessed Feb. 19, 2021), <http://www.originproject.info/about>.

¹² These could be SBIRs, OTAs, or other modern vehicles with minimal red tape. Recently published reports on countering malign influence have issued wide-ranging recommendations including: deploying special operations forces to areas “vulnerable to political warfare,” building “rapid-reaction information cells to track and counter” malign influence operations, and promoting civil society to “combine the values of accurate media with engagement skills and an understanding of how propagandists prey on polarization, inflaming divides.” These recommendations are already being put into action by Special Operations Command in Africa, using commercially available services to combat and attribute malign information operations about COVID-19 and other issues on the continent. The General Services Administration has awarded IST Research a Phase III SBIR contract to “support operations in the information environment for the special operations and Federal Government community.” Additionally, the U.S. Air Force and U.S. Special Operations Command have contracted with Primer to “automatically identify and assess suspected disinformation” using ML technology. See David Ronfeldt & John Arquilla, *Whose Story Wins: Rise of the Noosphere, Noopolitik, and Information-Age Statecraft*, RAND at 72 (July 2020), https://www.rand.org/content/dam/rand/pubs/perspectives/PEA200/PEA237-1/RAND_PEA237-1.pdf (citing or quoting other experts or reports); Dave Nyczepir, *SOCOM Looks to Combat Disinformation in Africa on New Governmentwide Contract*, FedScoop (July 27, 2020), <https://www.fedscoop.com/socofrica-disinformation-ist-research/>; *IST Research Awarded Five-year, \$66 Million GSA Contract*, IST Research (July 23, 2020), <http://www.istresearch.com/>.

globenewswire.com/news-release/2020/07/23/2066650/0/en/IST-Research-Awarded-Five-year-66-Million-GSA-Contract.html; *SOCOM and US Air Force Enlist Primer to Combat Disinformation*, Cision PR Newswire (Oct. 1, 2020), <https://www.prnewswire.com/news-releases/socom-and-us-air-force-enlist-primer-to-combat-disinformation-301143716.html>.

¹³ Nicholas Duran, et al., *2018 Webroot Threat Report*, Webroot (2018), https://www-cdn.webroot.com/9315/2354/6488/2018-Webroot-Threat-Report_US-ONLINE.pdf; Gary J. Saavedra, et al., *A Review of Machine Learning Applications in Fuzzing*, arXiv (Oct. 9, 2019), <https://arxiv.org/pdf/1906.11133.pdf>; Isao Takaesu, *Machine Learning Security: Deep Exploit*, GitHub (Aug. 29, 2019), https://github.com/13o-bbr-bbq/machine_learning_security/tree/master/DeepExploit; Catherine Stupp, *Fraudsters Used AI to Mimic CEO's Voice in Unusual Cybercrime Case*, Wall Street Journal (Aug. 30, 2019), <https://www.wsj.com/articles/fraudsters-use-ai-to-mimic-ceos-voice-in-unusual-cybercrime-case-11567157402>; *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine (2019), <https://doi.org/10.17226/25488>; Ben Buchanan, et al., *Automating Cyber Attacks: Hype and Reality*, Center for Security and Emerging Technology (Nov. 2020), <https://cset.georgetown.edu/research/automating-cyber-attacks/>; Nektaria Kaloudi & Jingyue Li, *The AI-Based Cyber Threat Landscape*, ACM Computing Surveys at 1-34 (Feb. 2020), <https://dl.acm.org/doi/abs/10.1145/3372823>; Dakota Cary & Daniel Cebul, *Destructive Cyber Operations and Machine Learning*, Center for Security and Emerging Technology at 5-23 (Nov. 2020), <https://cset.georgetown.edu/research/destructive-cyber-operations-and-machine-learning/>.

¹⁴ Max Heinemeyer, *Dissecting the SolarWinds Hack without the Use of Signatures*, Darktrace (Jan. 7, 2021), <http://www.darktrace.com/en/blog/dissecting-the-solar-winds-hack-without-the-use-of-signatures/>.

¹⁵ See *EINSTEIN*, U.S. Cybersecurity & Infrastructure Security Agency (last accessed Feb. 8, 2021), <https://www.cisa.gov/einstein>.

¹⁶ Gregory C. Wilshusen, *DHS Needs to Enhance Efforts to Improve and Promote the Security of Federal and Private-Sector Networks*. Government Accountability Office, (April 24, 2018), <http://www.gao.gov/assets/700/691439.pdf>. See also *Information Security: DHS Needs to Enhance Capabilities, Improve Planning, and Support Greater Adoption of Its National Cybersecurity Protection System*, Government Accountability Office (Jan. 28, 2016), <https://www.gao.gov/assets/680/674829.pdf>.

¹⁷ *President's FY 2021 Budget Request to Congress for the National Institute of Standards and Technology*, National Institute of Standards and Technology (2020), <http://www.nist.gov/system/files/documents/2020/02/11/FY2021-NIST-Budget-Book.pdf>.

¹⁸ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine at 17-19 (2019), <https://doi.org/10.17226/25488>.

¹⁹ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine at 12-19 (2019), <https://doi.org/10.17226/25488>.

²⁰ *Implications of Artificial Intelligence for Cybersecurity: Proceedings of a Workshop*, National Academies of Sciences, Engineering, and Medicine at 15 (2019), <https://doi.org/10.17226/25488>.

²¹ *Report*, U.S. Cyberspace Solarium Commission (March 2020), <https://www.solarium.gov/report>. [hereinafter CSC Report]

²² This recommendation modifies an existing Cyberspace Solarium Commission recommendation in order to reduce the risk of creating perverse incentives to avoid enforcement. See recommendation 4.2 and 4.4.4, CSC Report at 76, 83.

²³ See recommendation 4.4.4, CSC Report at 83.

²⁴ See recommendation 4.4.3, CSC Report at 82.

²⁵ See recommendation 4.5.3, CSC Report at 87.

²⁶ S. 3288, International Cybercrime Prevention Act of 2018, 115th Congress (2018), <https://www.congress.gov/bill/115th-congress/senate-bill/3288/text>.

Blueprint for Action: Chapter 1 - Endnotes

²⁷ *Report of the Attorney General's Cyber Digital Task Force*, U.S. Department of Justice at 124 (July 2, 2018), <https://www.justice.gov/archives/ag/page/file/1076696/download>.

²⁸ See recommendation 6.1.2, CSC Report at 114.

²⁹ See recommendation 1.1, CSC Report at 32.

³⁰ See recommendations 1.1.1 and 1.1.2, CSC Report at 32.

³¹ See recommendation 5.4, CSC Report at 87.

³² See recommendation 4.6, CSC Report at 88.

³³ See recommendation 4.3, CSC Report at 78.

³⁴ CSC Report, 78.

³⁵ See recommendation 4.1.1, CSC Report at 75.

³⁶ See recommendation 5.2, CSC Report at 101.

³⁷ The President's National Infrastructure Advisory Council detailed a similar recommendation to make cyber intelligence more actionable. *Transforming the U.S. Cyber Threat Partnership*, President's National Infrastructure Advisory Council at 8 (Dec. 12, 2019), <https://www.cisa.gov/sites/default/files/publications/NIAC-Working-Group-Report-DRAFT-508.pdf>.

³⁸ CSC Report at 102

Chapter 2: Foundations of Future Defense

Blueprint for Action

The Department of Defense (DoD) lags far behind the commercial sector in integrating new and disruptive technologies such as Artificial Intelligence (AI) into its operations. Technical, bureaucratic, and cultural challenges must be overcome to adopt AI to maintain the U.S. military advantage. By 2025, the DoD must put in place the foundations for widespread AI adoption, by: 1) Building the technical backbone; 2) Training and educating warfighters; 3) Accelerating adoption of existing digital technologies; 4) Democratizing development of AI; and 4) Investing in next-generation capabilities. To the maximum extent possible, these efforts should be coordinated with the Intelligence Community (IC) and other partners across the national security community.¹

Recommendation

Recommendation: Drive Change through Top-Down Leadership.

Maintaining the defense advantage in an AI-enabled future will require top-down leadership to overcome organizational barriers and create strategic change. Critically, civilian and military leaders across the DoD and the IC must coordinate more closely, aligning priorities, resources, and policies to speed technology adoption and research breakthroughs.

Actions for the Department of Defense and the Office of the Director of National Intelligence:

- **Establish a Steering Committee on Emerging Technology, tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.**²
 - o The Secretary of Defense and Director of National Intelligence should issue a directive immediately establishing the senior oversight committee listed above.
 - o The Steering Committee on Emerging Technology provides a forum to drive change, focus, and action on emerging technology that otherwise would not be prioritized. It will enhance intelligence analysis related to emerging technology; connect strategic vision to organizational change; focus concept and capability development on emerging threats; guide defense investments that ensure America's strategic advantage against near-peer competitors; and provide the authority to drive technology adoption and application by the Department.
- **Assign the tri-chaired Steering Committee on Emerging Technology responsibility for overseeing the development of a Technology Annex to the next National Defense Strategy.**³

Action for Congress:

- **In the National Defense Authorization Act (NDAA) for FY2022, establish a Steering Committee on Emerging Technology and National Security Threats and designate that it be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.**

Recommendation: Build the Technical Backbone

Recommendation

Integration of AI into DoD operations requires urgent investment in a modern digital ecosystem that will enable ubiquitous development and fielding at all levels—from the headquarters to the tactical edge. It is essential to establish a technical foundation that:

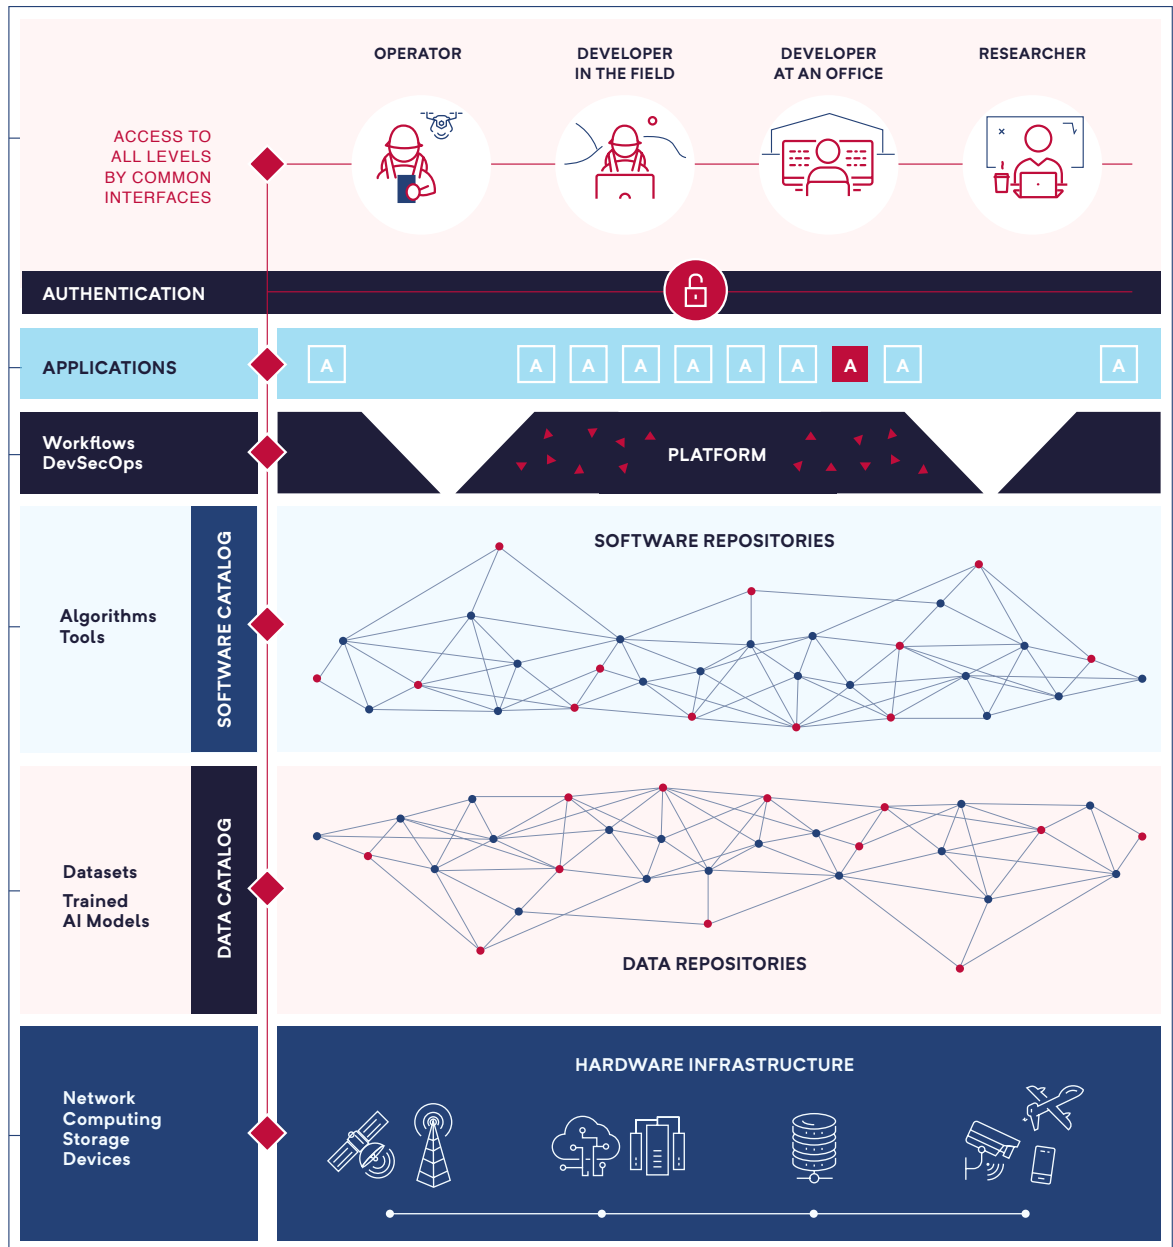
- 1) Provides access to leading cloud technologies and services for scalable computing;
- 2) Enables the sharing of data, software, and capabilities through well-documented and hardened application programming interfaces (API) with proper access controls; and
- 3) Gives all DoD developers and scientists access to the tools and resources they need to drive new AI capabilities. To this end, the figure below depicts the ecosystem managed as a multilayer stack of services, accessed through common interfaces and providing shared access to essential AI building blocks of data, algorithms, tools, trained AI models, and compute. This should be realized through a federated approach, building on existing resources and pathfinder efforts.⁴

The key elements that comprise the envisioned AI digital ecosystem are:

- *Contributors and Users.* A diverse, distributed network that includes development teams working at the tactical edge and at headquarters levels; private sector partners contributing trained models and applications; academic researchers working on open challenge problems; researchers working within a DoD lab; or international allies or partners co-developing interoperable AI capabilities.
- *Common Interfaces.* A service-oriented architecture where resources at each level of the stack are accessed and maintained through common APIs based on industry-standard protocols.
- *Authentication.* Enhancing both the sharing and the safeguarding of resources through a uniform policy and practice for managing authoritative, shared user attributes across classification levels to control who will build, use, or share AI building blocks.⁵
- *Applications.* Discoverable and accessible AI solutions ready for fielding through provisioned platform environments.⁶
- *Platforms.* Environments that support development, testing, fielding, and continuous updating of applications to diverse sets of contributors and users.⁷ These platforms include workflows and processes supporting the DevSecOps⁸ life cycle, MLOps⁹ for machine learning pipelines, and digital engineering.¹⁰
- *Software.* Federated software architecture¹¹ linking distributed repositories hosted across the Department by mission components, their software factories, and service labs, making software discoverable through a catalog.¹² Includes AI algorithms, data analysis tools, and tools supporting test and evaluation, verification and validation (TEVV)¹³ as well as processes and tools to support continuous Authorization to Operate (ATO) frameworks and reciprocity.¹⁴

- **Data.** Federated and secured data architecture linking distributed repositories across the department hosted by mission components, service labs, and enterprise programs, making data discoverable through a catalog.¹⁵ With appropriate access controls, this will facilitate finding, accessing, and moving desired data across the Department¹⁶ including data sets, associated data models, and trained AI models along with supporting documentation.¹⁷
- **Hardware Infrastructure.** Networking and communications backbone to transport ecosystem resources, particularly data, and provide seamless access and interchange between cloud computing and storage services.

AI Digital Ecosystem.



To accelerate the process of building on existing resources and pathfinder efforts, and to increase interoperability in the short term, DoD should determine a governance structure and develop necessary policies and guidance, draft a reference design, and make the technical investments in the network and in platform environments. Implemented correctly, the digital ecosystem will ensure force-multiplying common access and interoperability. The Blueprint for Action framework outlined below marries top-down coordination and direction with bottom-up mission implementation to realize an enterprise-wide ecosystem in a manner that does not slow or stymie innovation, but rather incorporates new capabilities at the speed of innovation and mission requirements.

Actions for the Department of Defense:

- **Establish Digital Ecosystem Leadership and Governance.**

- o The Secretary of Defense should direct the establishment of an enterprise-wide digital ecosystem to support capability development to maintain the technological superiority of the United States military.
 - To ensure senior leader oversight and sustained resourcing, the Secretary should assign the Steering Committee on Emerging Technology the responsibility to oversee the implementation and sustainment of the ecosystem.
 - The Secretary should assign the DoD Chief Information Officer (CIO) as the Executive Agent responsible for the ecosystem design, development, and operation.
- o The Steering Committee on Emerging Technology, coordinating with the DoD CIO, DoD Comptroller, Director of Cost Assessment and Program Evaluation, and appropriate acquisition and programming representatives from the military services, should produce a funding plan¹⁸ for the ecosystem within 90 days of the Secretary's direction.
- o The DoD CIO should form and chair an enduring digital ecosystem implementation working group¹⁹ to establish and maintain an open architecture, an evolving reference design, governance structure, and processes to include management and authorization for ecosystem functions and growth. The Steering Committee on Emerging Technology will ensure strategic direction and coordination, and pathfinder organizations will provide bottom-up and mission-oriented implementation.²⁰
 - The working group should report to the Steering Committee on Emerging Technology, add members when appropriate, and include representatives from:²¹
 - The Office of the DoD Chief Data Officer (CDO)
 - Component CIOs and CDOs
 - The Joint Artificial Intelligence Center (JAIC)
 - The Office of the Under Secretary of Defense for Research & Engineering (OUSD (R&E))

- The Office of the Under Secretary of Defense for Acquisition & Sustainment (OUSD (A&S))
 - The Office of the Under Secretary of Defense for Intelligence & Security (OUSD (I&S))
 - Service Acquisition Executives
 - The Office of the Director of Operational Test and Evaluation (DOT&E)
 - The Director of the Defense Advanced Research Projects Agency (DARPA)
 - Digital ecosystem pathfinders, including but not limited to, the Air Force's PlatformOne, Kessel Run, Space CAMP, the Navy's Black Pearl, the Army's CReATE, ADVANA, and the Army Futures Command Software Factory²²
- **Develop and Mandate Participation in a Digital Ecosystem Governed by an Open Architecture and Reference Design.**
 - o Within 12 months of the Secretary's direction to establish the ecosystem, the DoD CIO should work with the implementation working group to develop and publish an open, interoperable architecture²³ built on common interfaces based on industry-standard protocols along with an evolving reference design.²⁴
 - The open architecture and reference design should be owned by the DoD CIO and reviewed quarterly and updated through the working group.
 - An unclassified version of the open architecture and reference design should be published publicly for commercial capability providers.
 - o The Secretary of Defense should issue a memorandum that requires all new joint and service programs to participate in the digital ecosystem and adhere to the open architecture.²⁵ This should include a requirement that all existing programs develop a plan to participate and become interoperable with the digital ecosystem wherever possible by 2025.
 - Through the Joint Requirements Oversight Council (JROC), the Vice Chairman of the Joint Chiefs of Staff (VCJCS) and USD (R&E)²⁶ should ensure that all joint and service programs designed to meet joint capability needs adhere to the digital ecosystem open architecture.²⁷
 - o The DoD CDO, acting in coordination with the DoD Data Council, should ensure that the Data Strategy Blueprints for Action Annex developed by each Component under the DoD Data Strategy adhere to the digital ecosystem open architecture.
 - o The USD (A&S) should update the guidance governing the formatting requirements for deliverable data in contracts to be well-documented,²⁸ "non-proprietary formats designed for interoperability."²⁹
 - o The Steering Committee on Emerging Technology should lead an effort with the IC to assess additional ways to accelerate implementation and leverage the digital ecosystem, including designating service providers to proliferate applications across the enterprise and make them available for integration into complex mission solutions.³⁰ Wherever possible, the digital ecosystem's open architecture should leverage and interoperate with proven solutions from the IC such as the Information Technology Environment recommended in Chapter 5 of this report.

- **Establish a Strategic Data Node.**

- o The DoD CDO should make it a priority to create a linked, large-scale, cloud-based data repository (i.e., a node within the digital ecosystem) adherent to the data service interfaces specified in the ecosystem's open architecture. This would be a critical step to enable distributed development efforts by providing AI development teams secure access to authoritative data from diverse mission sets and functional areas, and serve as a prototype for the digital ecosystem reference design.³¹
 - The CDO should create this strategic data node by integrating digital ecosystem interoperability into the DoD ADVANA system³² and prioritize construction of enterprise data sets as recommended below.

- **Expand the Network and Communications Backbone to the Digital Ecosystem.**

- o The Department should fully fund its network and communications modernization effort as outlined in the DoD Digital Modernization Strategy,³³ require the DoD CIO to factor this into their list of highest priorities, and hold the DoD CIO accountable for expediting critical upgrades.

- **Create a Marketplace to Promote Democratization of AI Building Blocks.**

- o The DoD CIO, in accordance with the digital ecosystem governance and reference design addressed above, should create an AI marketplace for strategic exchanges of the essential AI building blocks.³⁴ The marketplace should include:
 - SoftEx — GitLab-like³⁵ software repository system³⁶ hosting AI algorithms, TEVV tools,³⁷ hardened AI software stacks, etc.
 - DataEx — a federated data repository system³⁸ of AI-ready data sets, documentation, and associated data models.³⁹
 - ModelEx — a federated repository system of trained models⁴⁰ generated from various types of AI approaches and techniques, including statistical machine learning.⁴¹
 - CloudEx — a cloud-agnostic, networked marketplace for pre-negotiated computing and storage services from a pool of vetted cloud providers.⁴²
- o Trusted partners (inside and outside government) should be able to develop solutions and products within secured environments of the ecosystem, offering monetized access to users.⁴³

- **Develop Prototypical Platform Environments within the Digital Ecosystem.**

- o The DoD CIO should work closely with the digital ecosystem pathfinder community to build a set of tailorable development environments for training AI systems using: data-driven statistical machine learning; the latest simulation and modeling capabilities to support reinforcement learning (e.g., digital twinning within an accurate world model); and complementary TEVV services.⁴⁴
- o The DoD CIO should work closely with the digital ecosystem pathfinder community to implement a set of prototypical platform environments⁴⁵ that support development, testing, fielding, and continuous update of AI-powered applications for diverse categories of contributors and users.⁴⁶

Action for Congress:

- **Prioritize funding for the Department's digital ecosystem and associated activities.**
 - o The Armed Services Committees should use the FY2022 NDAA to direct the Department to develop a resourcing plan for the digital ecosystem that establishes, sustains, and incentivizes use of its various components as enterprise-wide, enduring resources.
 - o The Committees should also authorize the obligation of funds to begin work on the ecosystem.

Recommendation

Recommendation: Train and Educate Warfighters

Warfighters need the following capabilities to effectively build and use AI-enabled systems:

- *Data-informed decision-making:* Data-informed decision-making uses data to generate insights and act on them. Data-driven organizations often make decisions more quickly, at lower levels in the organization, and with a stronger empirical foundation than organizations that rely primarily on intuitive or experience-based decision-making.⁴⁷
- *Computational thinking:* Service members need to better understand how to use information processing agents to perform beneficial calculations that could not be done quickly or efficiently by a human, rather than just representing human thinking in a digital format.
- *Maker culture:* Service members of all ranks and occupations need regular contact with AI-enabled machines, and should be able and encouraged to experiment with and participate in the development of new tools.
- *Human-machine teaming:* Military leaders need to understand how to effectively provide input to machines, interpret machine outputs, and critically, when to trust or not trust machine outputs.⁴⁸
- *Organizational transformation:* Leaders need to understand when and how to integrate AI-related tasks into their organization's priorities, allocate resources needed to build and maintain the AI stack, oversee the deployment and scaling of new systems, and how to effectively interact with and support the careers of their technical experts.

Component 1: Integrate Digital Skill Sets and Computational Thinking into Military Junior Leader Education.

Military junior leaders need to understand enough about AI to manage and operate AI-enabled organizations responsibly and effectively. Commanding and leading AI-driven systems and humans are very different fields. Leadership is even more complex in organizations that combine human and AI elements. The below skill sets will equip junior leaders with the fundamental skills needed.

Problem Definition and Curation. Military leaders need to understand problem curation, or the process of discovering the causal mechanisms that lead to problems, associated issues, stakeholders, and potential minimum viable products.⁴⁹ Poor problem definition

and curation can lead to projects that attempt to solve the incorrect problem, wasting significant amounts of time and money. This is particularly true for AI. Not all problems can be solved using the type of probabilistic reasoning performed by many algorithms, or with limited data sets. Also, many problems with potential AI solutions can be solved with much easier, less-resource-intensive techniques. Military leaders that understand problem curation will be better able to identify problems with potential AI solutions, and, just as importantly, problems that AI will not help solve. This would not only help with the use of AI but would also make junior leaders generally more productive.

A Conceptual Understanding of the AI Stack. The AI stack is a model that “provides a streamlined approach to visualize, plan, and prioritize strategic investments in commercial technologies and transformational research to leverage and continuously advance AI across operational domains, and achieve asymmetric capability through human augmentation and autonomous systems.”⁵⁰ A conceptual understanding of the AI stack would reinforce the importance of building structural solutions to data collection, management, curation, installation of sensors, and other underappreciated topics, and reduce attempts to add AI at the end of a project. It will also help military leaders better understand what part of their adversaries’ AI to target to degrade its effectiveness.

Data Collection and Management. Junior leaders need to understand how to collect and manage data and to use systems that do the same in a manner that prepares it for exploitation, and to operate in an environment where adversaries are doing the same. They also need to understand the causes, effects, and ethical implications of data bias. Training junior leaders to collect and manage data with the same degree of responsibility and expertise that they use for medical care and equipment maintenance would accelerate the government’s ability to create AI solutions, and to employ data-informed decision-making.

Understanding Probabilistic Reasoning and Data Visualization. Junior leaders need to know enough about probabilistic reasoning and data visualization to understand the outputs of their AI systems and their implications for a particular situation or environment. This is critically linked to understanding when to trust and not trust a system’s outputs, and other aspects of commanding and leading AI-driven systems. Notably, this does not require leaders to perform computational statistics, just to understand their output, a much less demanding task.

Data-informed Decision-making. To make data-informed decisions, leaders need to understand system thinking and critical thinking. System thinking combines all of the above to create an empirical but incomplete understanding of factors influencing a decision, and how both their system affects their AI and how their decision will affect their system. Critical thinking will help leaders understand the limits of AI, and the limits of data-informed decision-making processes that are based on imperfect information. This report references data-informed rather than data-driven decision-making because military leaders should never be bound by the imperfect information in front of them. Their critical

thinking, judgment, and intuitive understanding of both their system and their environment will always have a critical role to play, even as it is informed by decision-making aids.

Action for Congress:

- **Require the military services to integrate digital skills and computational thinking into pre-commissioning and entry-level training.**
 - o The Armed Services Committees should use the FY2022 NDAA to require the military services to integrate understanding problem curation, the AI life cycle, data collection and management, probabilistic reasoning and data visualization, and data-informed decision-making into existing, pre-commissioning or entry-level training for junior officers and training for non-commissioned officers within one year of the passage of the legislation.

Action for the Military Services:

- **Integrate digital skills and computational thinking into pre-commissioning and entry-level training.**
 - o The military services need to integrate understanding problem curation, the AI life cycle, data collection and management, probabilistic reasoning and data visualization, and data-informed decision-making into pre-commissioning or entry-level training for junior officers and training for both junior and senior non-commissioned officers. The military services can accomplish this by creating new modules or courses, or by integrating this training into existing training and education for commissioned and non-commissioned officers. Whenever possible, this training should include the use of existing AI-enabled systems and tools.

Component 2: Integrate Emerging and Disruptive Technologies into Service-level Professional Military Education.

While it is critical that military junior leaders better understand digital technology, military leaders must also understand how technology will affect warfare, their roles in their organizations, and how they should integrate new technology, both as they increase in rank and responsibility and as technology changes.

Action for Congress:

- **Require the military services to integrate emerging and disruptive technologies into service-level Professional Military Education.**
 - o The Armed Services Committees should use the FY2022 NDAA to direct the DoD to require emerging and disruptive technologies courses for officers within one year. The Armed Services Committees should also require the DoD to develop a training plan that incrementally builds the necessary skills in its officer corps.

Action for the Military Services:

- **Integrate emerging and disruptive technologies into service-level Professional Military Education.**

- o Course materials should address AI and other militarily significant emerging technologies, as identified by the military services and the USD (R&E), in coordination with the national laboratories, Federally Funded Research and Development Centers (FFRDCs), and University Affiliated Research Centers (UARCs).
- o Course materials should include an introduction to the latest technology, the benefits and challenges of adapting new technologies, how organizations successfully and unsuccessfully adopt these technologies, and ethical issues surrounding the uses of emerging technologies, including the impact of biases in these technologies.
- o As officers progress in rank, such courses should increasingly build the knowledge base, vocabulary, and skills necessary to better understand new threats/challenges, develop operational and organizational concepts, and incorporate technology into operations and operational support.
- o Military services should establish a mechanism that audits these courses annually to ensure that emerging technologies have been properly identified, and that the nomenclature, lexicon, definitions, and course content match changes in emerging technologies.

Component 3: Create Emerging and Disruptive Technology Coded Billets in the Department of Defense.

It is crucial that the DoD incentivize and increase the skill needed to introduce and field emerging and disruptive technologies within the military officer corps. The joint qualification process can serve as a model. The DoD already designates that certain critical billets must be filled by Joint Qualified Officers⁵¹ and different levels of joint qualification.⁵² To do this, the DoD should create emerging and disruptive technology designated billets for officers that require an emerging and disruptive technology qualification prior to assignment and a process for military leaders to become emerging and disruptive technology qualified. Emerging and disruptive technology qualified officers would add value in a number of areas for the services, including: 1) Assisting with acquisition of emerging technology; 2) Helping integrate technology into field units; 3) Developing organizational and operational concepts; and 4) Developing training and education plans.

Action for Congress:

- **Require the Department of Defense to create emerging and disruptive technology critical billets that must be filled by emerging technology certified leaders.**

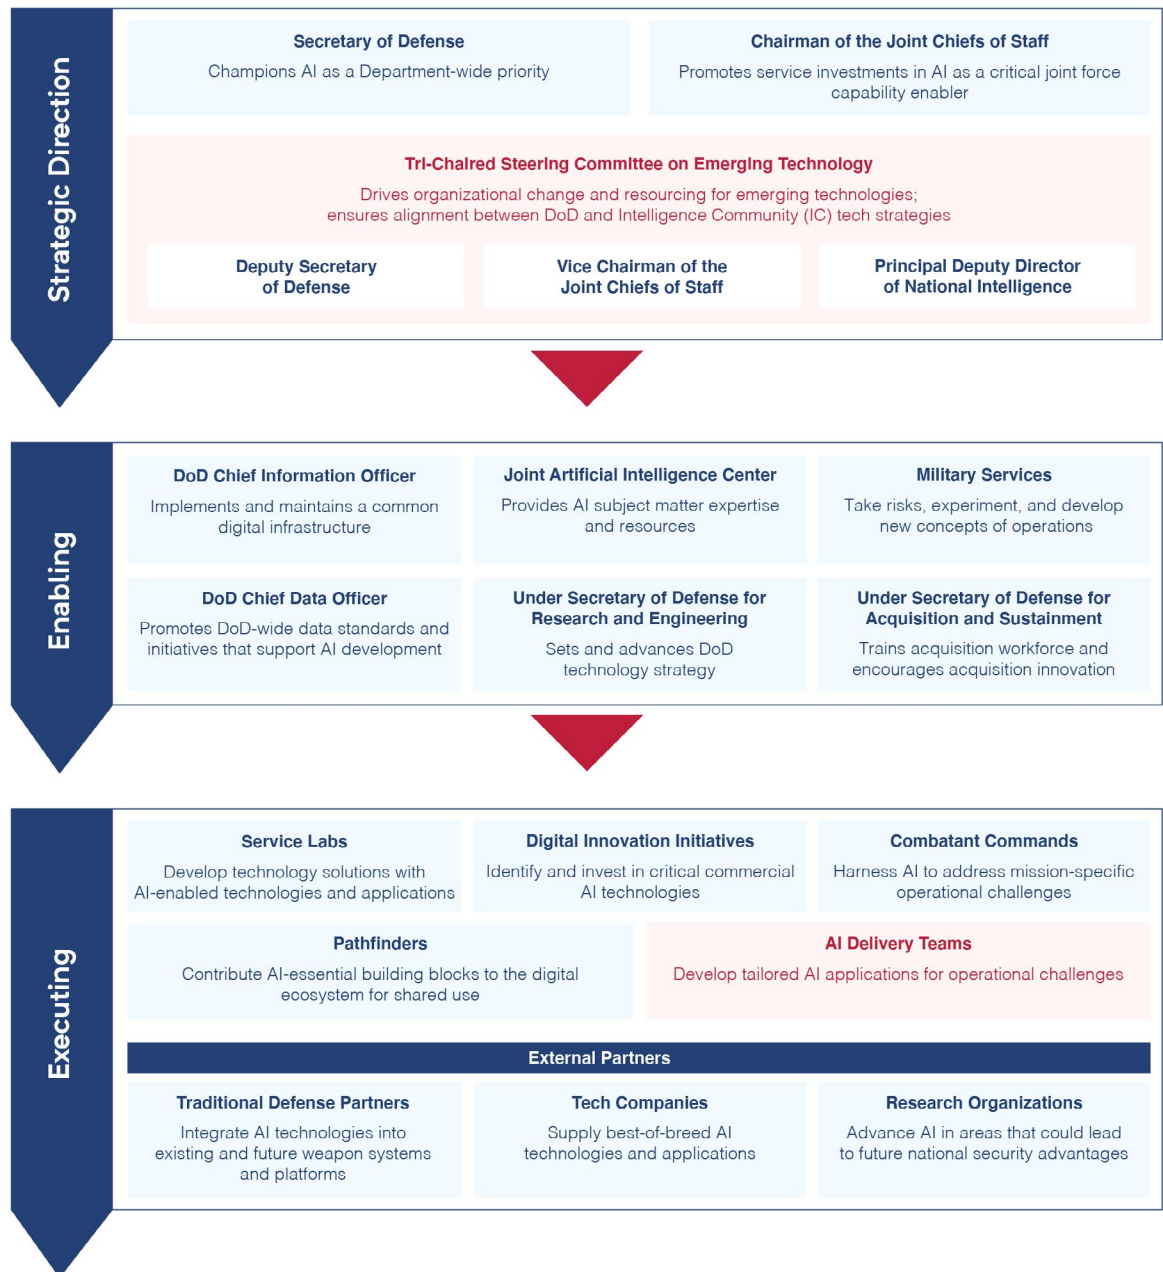
Actions for the Department of Defense:

- **Create billets that require officers to become emerging and disruptive technology certified before serving in the positions.**
 - o The Office of the USD (R&E) should define emerging and disruptive technologies.
 - o Billets include, but are not limited to, positions that develop military doctrine and/or operating concepts; positions within Force Structure, Resources, and Assessment

directorates; positions within Force Development directorates; and leadership positions at the operational and tactical levels within the military services.

- **Create a process for officers to become emerging and disruptive technology certified.**
 - o The process to become emerging tech certified would resemble the joint qualification system.
 - o Officers should become emerging technology qualified by serving in emerging technology focused fellowships,⁵³ emerging technology focused talent exchanges, emerging technology focused positions within government, and completing educational courses focused on emerging technologies.

Organizing DoD for AI Adoption.



Recommendation: Accelerate Adoption of Existing Digital Technologies

The Department must have an integrated approach to AI and other emerging technologies that ensures the U.S. military can continuously identify, source, field, and update capabilities faster than our competitors. This requires more targeted investment in dual-use technologies, ensuring system adaptability through a more agile budget and oversight process, and streamlining the acquisition process to shed those rules and regulations whose benefits are outweighed by the burdens imposed on the system. Critically, the Defense Acquisition System must shift away from a one-size-fits-all approach to measuring value from the acquisition process. Adherence to cost, schedule, and performance baselines is rarely a proxy for value delivered, but is particularly unsuited for measuring and incentivizing the iterative approaches inherent in AI and other software-based digital technologies. Unless the requirements, budgeting, and acquisition processes are aligned to permit faster and more targeted execution, the U.S. will fail to stay ahead of potential adversaries.

Component 1: Adopt Proven Commercial AI Applications for Core Business Processes.

Commercial AI applications for business processes can generate labor and cost savings, speed administrative actions, and inform decision-making with superior insights. To realize these benefits, DoD should initiate the digital transformation of its core administrative functions.

Efforts to apply business AI depend on the availability of clean, organized data. Significant resources are required to access, clean, and label enterprise data from the range of legacy business platforms.

DoD should create opportunities for bottom-up identification of AI use cases by incentivizing DoD organizations to deploy proven commercial applications tailored to their specific mission needs. Promising categories of commercial AI include: 1) Knowledge management applications such as intelligent search tools that index, retrieve, and display an agency's digital information, as well as collective intelligence and coaching tools that accumulate and exchange tacit knowledge across an agency's workforce; 2) AI-enabled tools that analyze business information to identify patterns, develop insights, and inform decision-making; and 3) Robotic Process Automation (RPA) tools including desktop assistants, bots, and other personal productivity applications that automate individual office functions.

Actions for the Department of Defense:

- **Prioritize construction of enterprise data sets across core DoD business administration areas.**⁵⁴
 - o The Deputy Secretary of Defense should:

- Assign the DoD CDO responsibility for working with institutional stakeholders to develop enterprise data sets for human resources, budget & finance, acquisition, logistics, retail, real estate, and health care.
 - Place special priority on the CDO building an enterprise data set that supports portfolio management of investments in emerging technologies, spanning budget requests, acquisition, contracting, and invoicing.⁵⁵
 - Assign the JAIC to support the DoD CDO in developing new methods for generating higher quality data for each core business administration area at the point of origin. This would include applying data tags that allow AI-enabled cross domain analyses.⁵⁶ As part of this effort, the JAIC should also look to develop or procure AI tools that continuously extract tagged information for analysis from enterprise data sets.
 - Ensure sufficient funding is included as part of the FY2023 budget request to provide data engineering services.
- The Secretary of Defense should issue a department-wide directive mandating the review and streamlining of policies and regulations wherever possible to increase and accelerate data sharing across agencies, with proper protections, building on the JAIC's Gamechanger AI prototype⁵⁷ to analyze and modernize the framework within which data access rules are enforced.
- **Launch a department-wide initiative to incentivize rapid deployment of commercial AI solutions for business functions.**
 - The Deputy Secretary of Defense should assign the JAIC, in its role as the Department's AI accelerator,⁵⁸ to administer allocation of matching funds, monitor and assess results, and disseminate best practices and lessons learned for the deployment of AI solutions for knowledge management, business analytics, and RPA across the Department, defense agencies, Services, and Combatant Commands.⁵⁹
 - The Secretary of Defense should issue a DoD directive mandating added flexibility and/or streamlining of policies and regulations wherever possible to increase and accelerate acquisition and deployment of commercial AI software, building on the JAIC's Gamechanger AI prototype to analyze and modernize the existing rules framework.⁶⁰
 - The Deputy Secretary of Defense should establish a \$100 million fund under the management of the JAIC to accelerate procurement and integration of commercial AI solutions for business applications. This would be used to provide matching incentive funds for agencies contracting with commercial AI vendors with approved solutions. The Deputy Secretary should also:
 - Consider leveraging the defense-wide review process detailed below to identify and reprogram sufficient funds to stand-up this fund by the beginning of FY2022.
 - Ensure sufficient funding is included as part of the FY2023 budget request to sustain the fund.

Action for Congress:

- **Congress should provide \$125 million as part of the FY2023 defense appropriations act for developing enterprise-wide data sets, and \$100 million for the fund to accelerate procurement and integration of commercial AI solutions for DoD business functions.**

Component 2: Network Digital Innovation Initiatives to Scale Impact.

Too often the Department's enthusiasm for innovation comes at the expense of impact and scale. Dozens of innovation offices across the Department and Services develop, transfer, and apply cutting-edge technology for national security uses.⁶¹ However, many of the initiatives that are focused on bridging the gap with the technology sector, to include AFWERX, NavalX, Army Applications Laboratory (AAL), and the Defense Innovation Unit (DIU), operate in silos and are limited in their ability to scale solutions. These pockets of successful bottom-up innovation have achieved some promising results, but disparate activities cannot translate to strategic change without top-down leadership to synchronize efforts and overcome organizational barriers.⁶² The Department should "network" programs that work to source cutting-edge technology solutions under the banner of "digital innovation initiatives" to execute a "go-to-market strategy" for digital technology that is supported at the highest levels of the Department.

Actions for the Department of Defense:

- **Designate an Executive Agent to coordinate the Department's digital innovation initiatives.**
 - o The Secretary of Defense should designate USD (R&E) as Executive Agent for the Department's digital innovation initiatives⁶³ and direct that USD (R&E) to coordinate closely with USD (A&S), DoD CIO, and DoD CDO to carry out the responsibilities associated with this role.⁶⁴
 - o As Executive Agent, USD (R&E) should facilitate access to resources,⁶⁵ provide strategic guidance, and offer other forms of institutional support to enable innovation organizations to execute their current mandates more effectively, without infringing on autonomy or inhibiting bottom-up experimentation.⁶⁶ USD (R&E) should work with the DoD CIO, the DoD CDO, and USD (A&S) as well as other institutional stakeholders as appropriate, to:
 - Develop a common digital platform for digital innovation initiatives to share data and procurement and development best practices,⁶⁷ track ongoing projects, connect with DoD program offices, and identify other means of collaboration.
 - Harness business AI tools to eliminate stovepipes and gain shared understanding of the digital innovation community, including investments and customers.⁶⁸
 - Identify and implement other metrics for the digital innovation initiatives to report as necessary, so long as they are lightweight and automated to the maximum extent possible.

- **Develop a “go-to-market” strategy for digital technology.**

- o USD (R&E) and USD (A&S) should issue a joint memo outlining a “go-to-market” strategy for digital technology to guide innovation organizations to pursue common objectives, based on the Technology Annex described below. This approach would coordinate efforts for effect and reduce duplication of effort, while preserving room for trial and experimentation with unexpected technologies or applications that could inform new operational concepts.⁶⁹
- o The Department should back this strategy with significant resources and top-down support. As described further in Chapter 11 of this report, DoD should set a target of increasing its contracts with early-stage technology firms by four times over the five-year Future Years Defense Program. To meet this goal, the Department should increase the procurement budgets of innovation organizations and other DoD entities to which innovation organizations hand off successful prototypes for production, as appropriate.
- o USD (R&E) should conduct annual investment portfolio reviews of digital innovation initiatives to assess alignment with strategy⁷⁰ and report findings to the Steering Committee on Emerging Technology.

- **Optimize operations to enable transition and scaling of AI solutions.**

- o USD (R&E), in partnership with USD (A&S), should assist innovation organizations in providing contracted vendors access and resources to build, deploy, and assure AI solutions often and at scale.⁷¹ In developing vendor contracts and planning customer journeys, digital innovation initiatives should consider the methods and means to:
 - Ensure that data access and data security requirements are included in contracts for AI systems that depend on data for training or operations.
 - Provide consistent access to end users as part of AI development processes and throughout the life cycle of the AI algorithm; and capture in contract terms.
 - Include AI testing and evaluation consideration as part of every development agreement.
 - Dedicate people and processes to onboard non-traditional vendors, migrate them onto the right networks and sandbox environments, and assist them in securing ATO.⁷²
 - Connect prototype contract recipients with DoD customers early in the technology development process and match program dollars with additional funding (SBIR, dedicated scaling funds, etc.) wherever possible.⁷³
 - Identify new opportunities for defense primes to team with non-traditional firms to adopt AI capabilities more quickly across existing platforms.⁷⁴
- o USD (R&E) should work with USD (A&S) to develop common reporting requirements to measure the impact of digital innovation initiatives, building off of ongoing efforts at DIU.⁷⁵ Collection of this data should be automated to the maximum extent possible and communicated routinely to Congressional defense committees. Reporting should consider:

- *Expansion of NSIB:* Number of awards made to companies with no previous DoD experience and percentage of these that receive follow-on contacts; or number of companies that receive recurring government revenue for first time and funding stability over consecutive quarters.
- *Rate of Transition:* Number of companies that receive follow-on production contracts.
- *Rate of Scaling:* Number of prototype contract recipients that reach \$10 million, \$50 million, \$100 million, \$500 million, and \$1 billion in total DoD contracts annually.
- *Reach of Products:* Number of users⁷⁶ that are benefiting from the product in one year, three years, 10 years, etc.
- *User Experience:* User feedback on the product (scale 1-10).
- *Company Acquisition Process Experience:* Company feedback on the new acquisition process (scale 1-10).
- *Operational/Enterprise Impact:* Actual or projected operational or fiscal return on investment (e.g., initiative addressed an operational gap; innovative RPA reduced production time or man-hours).

Component 3: Expand Use of Specialized Acquisition Pathways and Contracting Approaches.

AI technologies are incompatible with the lengthy, linear processes typical of traditional DoD capabilities acquisition.⁷⁷ Recent policy reforms such as the rollout of the Adaptive Acquisition Framework⁷⁸ (AAF) and associated resources—such as the Contracting Cone⁷⁹—are positive steps to move the Department away from a “one-size-fits-all” approach to acquisition. However, use of the specialized pathways and authorities⁸⁰ within the Framework is inconsistent and disincentivized.⁸¹ The traditional acquisition process remains the default for most acquisition professionals—many of whom are neither incentivized nor properly equipped to make use of the full resources at their disposal through the Framework.

To accelerate delivery of AI-enabled technologies to the warfighter and increase their operational relevance, DoD must build the capacity to use the full breadth of acquisition pathways and contracting approaches.⁸² Acquisition professionals must have a sufficient understanding of digital and emerging technologies in order to thoughtfully apply these tools. Given the speed of advancements in AI and other software-based technologies, this requires a shift to a continuous learning mindset and a different approach to training for acquisition professionals in which the target metric for success is not course completion, but rather the ability to apply what is learned and impact mission outcomes. DoD should coordinate acquisition workforce training initiatives relative to digital and emerging technologies ongoing across the Department and continuously assess acquisition workforce capability needs. Importantly, the DoD must also ensure acquisition personnel have common access to available digital technology courses and best practices as well

as a community of experts that illustrate how specialized authorities can be used to deliver best of breed technologies.

Actions for the Department of Defense:

- **Accelerate training of acquisition professionals and senior leaders on the AAF, Contracting Cone, and Digital Technologies.**
 - o The Secretary of Defense should develop a set of best practices in the use of new acquisition pathways⁸³ and direct USD (A&S) and Component Acquisition Executives to train the right acquisition professionals and DoD senior leaders and executives on the AAF, the Contracting Cone, and best practices for the use of these flexibilities, within one year.
 - o USD (A&S) should also work closely with USD (R&E), the JAIC, the Service Acquisition Executives, and the Component Acquisition Executives to implement a coordinated approach to training acquisition professionals and senior leaders on cross-functional specialties relative to emerging technologies. The approach should amplify and harmonize ongoing workforce training efforts⁸⁴ related to AI, data analytics, software, and digital engineering and look to leverage training or courses that can be procured off-the-shelf or as a service.
- **Leverage public-private talent exchanges to infuse technical expertise into the acquisition corps.⁸⁵**
 - o The Secretary of Defense should direct that acquisition professionals are considered among the highest priority to participate in public-private talent exchanges.⁸⁶
- **Establish enterprise learning platforms, course catalogs, and knowledge management tools for acquisition personnel and make them available Department-wide.⁸⁷**
 - o USD (A&S) should invest in and scale appropriate learning platforms, course catalogs, and knowledge management tools and create incentives for their use by FY2022. These resources should catalog available training⁸⁸ and best practices⁸⁹ and make relevant experts and specialists discoverable for acquisition professionals Department-wide.
- **Continuously assess existing acquisition workforce capabilities and evolve training for acquisition professionals.**
 - o The Secretary of Defense should direct that USD (A&S) work with the Service Acquisition Executives, Component Acquisition Executives, USD (R&E), and the JAIC to ensure curricula and approach to training⁹⁰ for acquisition professionals is consistently and appropriately updated to support the Technology Annex to the National Defense Strategy, as described below.

Action for Congress:

- **Authorize the use of a rapid contracting mechanism for the software acquisition pathway.**
 - o The Armed Services Committees should direct the Secretary of Defense to develop a rapid contracting mechanism to support the AAF's software acquisition pathway.⁹¹ The mechanism should include:

- A value-based price evaluation model.
- An independent, non-advocate cost estimate developed in parallel with engineering and leveraging agile cost estimation best practices.
- Performance metrics intended to measure value that can be automatically generated by users and shared as requested by DoD officials and congressional defense committees.

Component 4: Modernize the Budget and Oversight Processes for Digital Technologies.

The DoD's budget process requires that funds be requested two years in advance of their execution and focuses planning within the five-year Future Years Defense Plan (FYDP). Resources are allocated to program elements (PEs) that are defined at the system level⁹² and based upon cost buildups for pre-determined and highly specified system requirements.⁹³ In addition, the life-cycle-phased appropriation categories⁹⁴ that govern the DoD budget structure run counter to the iterative process inherent to AI and other software-based technologies.⁹⁵

This construct creates a paradigm unfriendly to the speed, adaptation, risk-taking, and joint force cohesion necessary to compete in an AI-enabled threat environment. Senior leaders champion the need to experiment⁹⁶ and “fail fast,” but the budget process prevents the allocation of funds without a justification clearly tied to program objectives. At the same time, the DoD requirements process—responsible for formulating the basis of those program objectives—assumes a linear and sequential relationship between requirements and technology.⁹⁷

To adapt faster than our adversaries, DoD must have a requirements and budget process that: 1) Prioritizes joint force capabilities and aligns resources accordingly; 2) Enables experimentation, iteration, and continuous development—especially for AI and digital technologies where persistent user feedback is critical; and 3) Balances speed, scale, and risk depending on the technology or capability being delivered.

Implementation of the large-scale institutional changes required to achieve this vision will take time and equal commitment from both DoD and Congress. In the near term, DoD and Congressional leaders should focus on generating mutual trust by establishing pilot programs to demonstrate the impact of reforms to the budget and requirements process relative to AI. The inclusion of support for the Department's Budget Activity 8 pilot program in the FY2021 defense authorization and appropriations acts represents positive progress to this effect.⁹⁸

Below are recommended steps that DoD and Congress should take immediately and over the longer term to create a modern budget and requirements process that supports the application of AI at speed and scale.

Immediate Actions for the Department of Defense:

- **Reorient the Joint Requirements Oversight Council (JROC) process to focus on Joint and Cross-Domain Capability.**
 - The Chairman of the Joint Chiefs of Staff should appoint the USD (R&E) Co-Chair and Chief Science Advisor of the JROC.⁹⁹
 - The Chairman of the Joint Chiefs of Staff should direct that the JROC charter be updated to reflect USD (R&E) as Co-Chair and Chief Science Advisor with responsibility for:
 - Delivering technology assessments and trend reports that inform JROC deliberations on future military requirements; and
 - Validating the technical feasibility of requirements developed by the services and ensuring they comply with the reference design for the digital ecosystem recommended above.
- **Make supplemental funding available to drive operational prototyping, scale, and transition of AI technologies.**
 - The Secretary of Defense should establish a dedicated AI fund as a pilot under the management of USD (R&E) to mature, operationally prototype, and transition exceptionally promising AI-enabled technologies that align with applications identified in the Technology Annex as described below. In doing so, the Secretary of Defense should direct:
 - Inclusion of approximately \$200 million for the fund in the FY2022 budget request.
 - USD (R&E), in collaboration with the JAIC and the military services, should establish clear metrics for success and a time horizon upon which to stand-up additional similar funds for specific technologies or capabilities.
- **Accelerate efforts¹⁰⁰ to implement a portfolio management approach for requirements and budget.**
 - The Deputy Secretary of Defense should produce a proposal for consideration in the FY2022 defense authorization and appropriation bills to establish a pilot to test a portfolio management approach for requirements and budgeting for at least one joint capability area, such as Command and Control, in FY2023.¹⁰¹ The proposal should:
 - Establish a reasonable ceiling value for the portfolio.
 - Consider consolidation of program elements and the creation of a Program Executive Office or other organizational entity empowered to resource and oversee programs designed to meet the joint capability need.
 - Request reprogramming authority to drive a “fail fast” mentality, promote experimentation and early prototyping, and quickly integrate new capabilities.
 - Provide recommendations on adjusted reporting guidance and justification documents, including metrics and mechanisms,¹⁰² that will allow Congress to conduct appropriate approval and oversight.

- o The Deputy Secretary of Defense should also produce a separate proposal to establish a pilot to test mission-focused budgeting and appropriations (e.g., a Mission Element). The proposal should be developed in coordination with a Combatant Command and organized around a high-priority operational challenge as identified by the Joint Staff. It should:
 - Consider more flexible funding mechanisms, including reprogramming authorities, applied across existing, relevant service programs to promote digital modernization and integration of AI technologies, interoperability, and new development or prototyping efforts for the specific operational challenge.
 - Provide recommendations on adjusted reporting guidance, and justification documents, including metrics and mechanisms, that will allow Congress to conduct appropriate approval and oversight.

Immediate Actions for Congress:

- **Update Title 10, Section 181 to designate USD (R&E) Co-Chair and Chief Science Advisor to the JROC.**
- **Direct the Secretary of Defense to establish the dedicated AI fund.**
 - o Congress should include a provision in the FY 2022 National Defense Authorization Act directing the establishment of an AI fund under USD (R&E) and appropriate at least \$200 million to support it as a pilot.¹⁰³
- **Support the continuation of the Budget Activity 8 pilot program in FY2022 and direct the Department to add an S&T project to the pilot programs.**
 - o Congress should continue to support the DoD software and digital technologies pilot program designed to allow for flexibility in funding the full life cycle of development, procurement, deployment, assurance, modifications, and continuous improvement for digital technologies.¹⁰⁴
 - o Congress should support DoD expanding the pilot in FY2022 to include a program that explicitly supports an S&T development effort in order to effectively test the impact of the single funding mechanism for the entirety of the AI life cycle, including early-stage research and development.

Longer-term Actions for the Department of Defense and Congress:


- **Establish a single appropriation and budget structure for software and digital technologies by FY2023.**
 - o Congress should build on the BA 8 pilot and establish a single appropriation for software and digital technologies that is exempt from the traditional programming or planning process and can be used as a single source of funding for the full life cycle of capability delivery and continuous engineering.
 - o The Department and Congress should collaborate to develop and implement a budget structure and transparent oversight process for the new software and digital technology appropriation that enables agile development of AI technologies and capability portfolio management.¹⁰⁵
- **Identify and implement successful portfolio- and mission-based budgeting constructs at scale across DoD.**

- o The Department and Congress should look to BA 8 as an example of how to apply a similar approach to monitoring and scaling portfolio- and mission-based budgeting. Based on metrics and oversight of the pilots over an appropriate timeline, DoD and Congress should determine what approaches to implement more broadly.

Recommendation*Recommendation: Democratize AI Development*

An AI-enabled threat environment requires our forces to be able to develop and deploy solutions nearly as quickly as threats arise. However, our forces frequently lack the infrastructure, tools, talent, and support to solve their challenges locally and with modern technology.¹⁰⁶ The JAIC cannot develop and proliferate AI applications for every user group or mission area within the DoD. To accelerate adoption of AI, the Department must create the technical infrastructure and organizational structures that pair top-down strategy with bottom-up development.

Democratize
Development.



Component 1: Leverage the JAIC as the Department's AI Accelerator.

The JAIC should serve as the Department's "AI accelerator" and central node for AI-related information. In this role, the JAIC would maintain critical situational awareness of AI stacks across the Department (i.e., options, including applications, available within the digital ecosystem that mission owners can leverage to enable local development efforts)

P

and provide the expertise and resources necessary to enable distributed development efforts.

Actions for the Department of Defense:

- **Designate the JAIC as the Department's AI Accelerator.**

- o The Deputy Secretary of Defense should issue a memorandum¹⁰⁷ designating the role of JAIC as the DoD enterprise's AI accelerator with responsibility for:
 - Developing tailorable AI applications to address high-level, cross-domain challenges and shared problems, and making them available through the digital ecosystem as enablers for development teams across the enterprise.
 - Administering a matching fund to incentivize integration of commercial AI solutions for business functions across the Department.
 - Collecting best practices (including best-of-breed AI applications) from industry, academia, and across the enterprise and making them accessible for the broader DoD developer community.¹⁰⁸
 - Providing AI subject matter expertise and assistance to DoD Components to inform strategy, policy, and technical approaches. This would include:
 - Participating as a member of the Steering Committee on Emerging Technology.
 - Contributing to the development of a reference design for the DoD AI digital ecosystem and associated governance policies.¹⁰⁹
 - Advising on integrating the appropriate governance frameworks for responsible use of AI into policies and procedures.¹¹⁰
 - Advising on TEVV policies and capabilities for AI.
 - In coordination with the Under Secretary of Defense for Policy, serving as the Department's lead for AI-related international engagement.
 - Developing a common AI TEVV framework,¹¹¹ in coordination with DOT&E and any other appropriate stakeholders, that integrates testing as a continuous part of requirements specification, development, deployment, training, and maintenance and includes run-time monitoring of operational behavior.¹¹²
 - Identifying, procuring, and orchestrating AI development tools and making them available through the digital ecosystem software exchange¹¹³ described above to enable distributed development efforts.¹¹⁴
 - Making available enterprise-wide contracting vehicles (e.g., Blanket Ordering or Purchase Agreements) for talent¹¹⁵ and AI technical services¹¹⁶ and continuously onboarding new companies.¹¹⁷
 - Coordinating with USD (R&E) on AI-related elements of the go-to-market strategy discussed above.
 - Integrating with nation-wide initiatives within other agencies and departments, as directed by the President.

- **Build technical support capability.**

- o The JAIC should grow and train a staff of resident experts¹¹⁸ that can provide support to users across the enterprise akin to an “AI help desk,” to include providing technical and policy consultation and advice; implementing solutions for small problems; and facilitating connections of support (for larger problems).¹¹⁹

Component 2: Embed AI development capabilities in support of operations.

The Department must ensure operators are paired with technologists at every echelon. Doing so will institutionalize user-centered; agile development, improve the speed and operational relevance of solutions delivered; and build trust and confidence in AI-enabled systems. Implementation of the actions below will create a networked support structure to enable bottom-up AI development extending from the tactical edge to the JAIC.¹²⁰

Actions for the Department of Defense:

- **Establish integrated AI delivery teams at every Combatant Command (CCMD).**

- o The Secretary of Defense should direct each Combatant Commander to stand-up an AI delivery team dedicated to developing and deploying AI applications to support operational units.¹²¹
- o Teams should be staffed with the appropriate talent to manage the full life cycle of AI solutions, including in disciplines such as data science, AI testing and model training, software engineering, product management, and full stack development.¹²² AI Delivery teams should be responsible for:¹²³
 - Finding, tailoring, and fielding applications from the digital ecosystem (e.g., those developed by other CCMDs, Service software factories, or the JAIC).
 - Developing additional sustainable mission applications as needed.
 - Contributing new and tailored applications to the digital ecosystem for use across the CCMD(s) to meet common challenges.

- **Integrate forward-deployed development teams with operational units.**

- o Each Combatant Commander should develop and implement a plan for the integration of forward-deployed development teams to act as the local customer interface for the AI delivery team with each operational unit.¹²⁴ Forward-deployed development teams should:
 - Work side-by-side with warfighters to identify problems and opportunities that could be met with AI applications.
 - Leverage the digital ecosystem to provision development environments and tools to produce “quick wins” to improve capabilities and generate efficiencies.¹²⁵

Recommendation

Recommendation: Invest in Next Generation Capabilities

The DoD must have an enduring process that clearly identifies, prioritizes, resources, and tracks¹²⁶ critical technologies over multiple time horizons. This will drive an investment strategy that pursues technology applications that close key capability gaps and optimize current operational concepts, and simultaneously makes bets on disruptive technologies to enable transformative capabilities and operational concepts over the long term.

Component 1: Increase investments in Science & Technology (S&T) and AI R&D.

To compete and win in AI-enabled warfare, the propagation of technology from core AI research to broad AI applications must expand drastically.¹²⁷ Across the board, increases in all lines of AI research¹²⁸ are called for, with particularly large increases in research funding required to advance key areas, such as developing methods for human-machine teaming and deploying trusted AI applications through rigorous methods for TEVV.

Action for the Department of Defense:


- **Commit to building budgets that invest at least 3.4%¹²⁹ of the annual DoD budget in S&T and allocate at least \$8 billion for research and development of core AI.**
 - o Particular focus should go toward strengthening the AI research budgets at organizations where AI expertise is centered, such as DARPA, the Office of Naval Research, the Air Force Office of Scientific Research, the Army Research Office, and the Service Laboratories.

Action for Congress:

- **Support DoD budget requests for amplified funding of AI R&D and AI-related initiatives.**

Enhanced AI R&D Investment, FY 2015–2030

Source: Govini and NSCAI



Enhanced AI R&D
Investment, FY
2015-2030.

Core AI spending is comprised of RDT&E programs to develop AI applications such as machine learning/deep learning, collaborative behavior, computer vision; human-machine teaming, automated reasoning, robotic autonomy, automated data fusion, and self-healing networks. Together with commercially sourced AI applications, core AI spending enables AI adoption. We define AI adoption as the sum of AI-enabled and AI-enabling investments (defined below).

AI-enabled programs develop (in the case of RDT&E programs) and field (in the case of procurement programs) the gamut of DoD warfighting and business systems, incorporating Core AI applications for analyzing, automating, communicating, maneuvering, monitoring, sensing, and many other tasks. While AI spending is usually a small percentage of these programs, their system's performance may be critically dependent upon the incorporation of core AI.

AI-enabling programs include technologies such as cloud computing and advanced microelectronics required to support the deployment of effective AI capabilities at scale.

Component 2: Retire Legacy Systems Ill-Equipped to Compete in AI-Enabled Warfare.

In the face of new budget realities, the Department must undergo an aggressive portfolio rebalance to ensure sustained room in its budget for emerging technologies like AI.¹³⁰ This will require DoD to make hard decisions on where to divest, and identify opportunities and timelines to upgrade or phase out legacy systems, as it continues to invest in new systems. However, the Department must also approach new systems differently. Rather than continuing to build large, monolithic platforms while competitors invest heavily in attritable systems, the DoD should focus on speed. DoD should drive investments into rapid prototyping and modular system design to develop and field new capabilities at a rate that allows U.S. forces to continuously out-adapt the adversary.

Actions for the Department of Defense:

- **Institutionalize an enduring defense-wide review and decision-making process,¹³¹ prioritized to the threat, to divest of legacy systems.**
 - o The Secretary of Defense should direct the Service Secretaries, USD (A&S), the Defense Agencies, and DoD Field Activities to evaluate the relevance and resiliency of all platforms and systems against emergent threats, and ruthlessly divest from systems and platforms deemed too costly or ineffective to equip with AI or make compatible with AI-enabled systems and architectures.¹³²
 - o The Service Secretaries and USD (A&S), in comparing the risk/reward tradeoffs between new versus old technologies and operating concepts, should leverage AI technologies as decision support tools.
 - o The Director of CAPE should enforce decisions to divest or reduce funding through the program review process.
 - o The Service Secretaries and USD (A&S) should explore options for updating legacy systems with leading-edge technologies to buy time for required long-term modernization projects.
- **Evaluate AI alternatives prior to funding new major defense acquisition programs.**
 - o The Secretary of Defense should issue a memorandum directing that all new major defense acquisition programs must conduct a thorough evaluation of AI alternatives as part of their analysis of alternatives prior to funding for major defense acquisition programs.¹³³
 - o USD (R&E) and the JAIC should provide support to program offices conducting such analysis by providing subject matter expertise informed by technology scouting and an awareness of the capabilities in the R&D pipelines across the S&T enterprise.

Action for Congress:

- **The Congressional defense committees should support the Department's hard decisions when presented with evidence that divestment or defunding can enable a more competitive force posture.**

Component 3: Create an integrated technical intelligence program¹³⁴ and a supporting community of practice.

To effectively leverage scientific and technological breakthroughs for competitive advantage, DoD must have a sophisticated technical intelligence program that monitors developments as they progress from basic research to prototype to fielded capabilities, understanding the R&D roadmaps of the private sector wherever possible. This intelligence must be global in scale, monitoring emerging technologies in near real time, especially in the rapidly evolving field of AI. The intelligence must be actionable, informing prioritization of resourcing and providing decision-makers the ability to continuously update technology roadmaps for our national security agencies.

Such a technical intelligence program should provide inputs to the proposed Technology Annex to the National Defense Strategy¹³⁵ in three main areas: 1) An understanding of the current and future threat capabilities in the R&D, production, and sustainment pipelines of our adversaries; 2) An understanding of the current and future friendly capabilities in the R&D, production, and sustainment pipelines of the U.S. government and allied partners; and 3) An understanding of emerging military and dual use technologies worldwide available for integration into national security capabilities.¹³⁶

Actions for the Department of Defense:

- **Transform the Strategic Intelligence Analysis Cell.**

- o USD (R&E) should reconceive the Strategic Intelligence Analysis Cell (SIAC)¹³⁷ as a robust analytic hub that marshals DoD, IC, and other technology scouting capabilities for strategic effect.¹³⁸ The SIAC Director should report directly to the USD (R&E).
- o SIAC should convene an interagency technology scouting community of practice from the service laboratories, OSD (including DARPA and DIU), innovation initiatives within the military services (such as AFWERX and AAL), the Departments of Energy and Homeland Security, university-affiliated research centers, federally funded research and development centers, CCMDs, and international security partners. This community of practice should:
 - Establish a federated approach to provide USD (R&E) with inputs to produce and continuously update the Technology Annex.
 - Conduct analytic exchanges and wargames to assess future technology scenarios and include AI to the maximum extent possible.¹³⁹
 - Develop rigorous technology forecasting capabilities, leveraging best practices from academia and the private sector.
 - Engage with industry and update requirements for technology scouting tools and data.
- o In order to leverage private industry more effectively, SIAC should maintain knowledge of private market investments relevant to the technologies and capabilities outlined in the Technology Annex.
- o In order to locate existing DoD capability gaps and potential solutions, SIAC must receive technical details at all levels of classification on current programs of record from OSD (A&S) and the armed service's acquisition executives, as well as technical details on RDT&E programs from OSD (R&E) and the technology scouting community of practice described above.
- o SIAC should establish a technology fellows program, inviting organizations in the technology scouting community to nominate personnel for short-term (three- to 12-month) assignments with SIAC where they would work side-by-side with SIAC analysts. This program should:
 - Build interdisciplinary teams of technologists and warfighters to conduct in-depth investigations of emerging technologies, initiating direct contacts with academia and industry in addition to passive data collection.

- Circulate personnel through the tech fellows' program into key roles in experimentation and concept development activities across OSD and the military services.
 - Develop personnel with greater understanding of emerging technologies across the national security community.
 - Leverage hiring authorities from the Public–Private Talent Exchange Program and the Intergovernmental Personnel Act to include fellows from industry, academia, and other government agencies to enhance access to non–DoD research and perspectives.¹⁴⁰
- o SIAC should acquire or develop research tools for use by the technology scouting community of practice, including AI-enabled analysis of large commercial databases, classified threat intelligence, and the technology investment portfolios of the United States Government and its allies.

Actions for Congress:

- **Congress should appropriate an additional \$10 million to USD (R&E)'s budget for the technology fellows program and AI-enabled technology scouting tools and data.**

Component 4: Develop a Technology Annex to the National Defense Strategy.

To identify where and how to direct scarce resources, the DoD should formulate its investment strategy as a classified Technology Annex to the National Defense Strategy (NDS) produced by the Department's Chief Technology Officer, USD (R&E). The Appendix should: 1) Identify emerging technologies and applications required to solve the operational challenges outlined in the NDS; and 2) Outline a clear plan for pursuing these technologies and applications. This plan should account for existing technologies, including dual-use commercial technologies, and drive rapid integration of these technologies to close near-term capability gaps.¹⁴¹ The plan should also help inform the agenda for DARPA and the DoD labs, by identifying disruptive technology elements and applications that warrant longer-term, exploratory investments. Finally, the plan must take into account industry's comparative advantage in available R&D capital and include a consistent and transparent approach to messaging defense technology priorities to build and broaden the industrial base.¹⁴²

Actions for the Department of Defense:

- **Develop a Technology Annex to the National Defense Strategy.**
 - o The Secretary of Defense, with support from the Director of National Intelligence, should issue a memo directing the Steering Committee on Emerging Technology to oversee the development of a comprehensive classified Technology Annex as a component of the next NDS and assign USD (R&E) as the Executive Agent responsible for producing the Technology Annex.
 - The Technology Annex should identify emerging technologies and applications that are critical to enabling specific capabilities for solving the operational challenges outlined in the NDS.

- o The Steering Committee on Emerging Technology described above should ensure that the Technology Annex sets clear guidance that drives prioritization and resourcing, while allowing enough flexibility for subordinate organizations to implement that guidance as best suits their mission. At a minimum, the Technology Annex should include:
 - Identified intelligence support requirements, including how the IC analyzes the global environment and monitors technological advancements, adversarial capability development, and emerging threats.
 - Identified functional requirements and technical capabilities necessary to enable concepts that address each challenge.
 - A prioritized, time-phased plan for developing or acquiring such technical capabilities that takes into account R&D timelines, a strategy for public private partnerships, and a strategy for connecting researchers to end users for early prototyping, experimentation, and iteration.
 - This should include roadmaps for designing, developing, fielding, and sustaining the technologies and applications to address the operational challenges outlined in the NDS.
 - These roadmaps should account for and leverage existing commercial-off-the-shelf/dual-use technologies and identify areas where defense-specific solutions are needed.
 - The roadmaps should use quantitative technological forecasting methods developed in academia and industry to identify disruptive technologies.
 - Identified additional or revised acquisition policies and workforce training requirements to enable DoD personnel to identify, procure, integrate, and operate the technologies necessary to address the operational challenges.
 - A prioritized, time-phased plan for integrating technology into existing DoD exercises that support the NDS.
 - Identified infrastructure requirements for developing and deploying technical capabilities, including data, compute, storage, and network needs; a resourced and prioritized plan for establishing such infrastructure; and an analysis of TEVV requirements to support prototyping and experimentation and a resourced plan to implement them.
 - Identified joint capability and interoperability requirements and a resourced and prioritized plan for implementation.
 - Consideration of human factor elements associated with priority technical capabilities, including user interface, human-machine teaming, and workflow integration.
 - Consideration of interoperability with allies and partners, including areas for sharing of data, tools, and operational concepts.
 - Flexibility to adapt and iterate appendix implementation at the speed of technological advancement.
- **Steward Implementation of the Technology Annex in Coordination with the Intelligence Community.**

- o The Secretary of Defense should direct the Steering Committee on Emerging Technology to steward implementation of the Technology Annex, to include coordination with the IC; and establish a reporting structure and metrics to monitor the implementation of each technology roadmap to ensure each effort is resourced properly and progressing sufficiently.
- o The Steering Committee on Emerging Technology should ensure common technical requirements are developed to align with the digital ecosystem's open architecture and are adhered to for the acquisition of emerging technologies identified in the Technology Annex. The standards should be coordinated across DoD and the IC.¹⁴³
- o The Steering Committee should conduct (at least) an annual review of the Appendix and ensure that guidance, policy, and implementation evolve at the pace of technological change.

Component 5: Clearly communicate defense technology priorities to industry.

DoD must leverage industry's comparative advantage in available R&D capital as part of its investment strategy. To do so effectively, the Department must adopt a consistent and transparent approach to messaging defense technology priorities that enables Defense primes and non-traditionals to plan and invest more to help meet DoD R&D needs, and further supports the Department's efforts to attract venture-backed companies.

Action for the Department of Defense:

- **Publish unclassified emerging technology R&D objectives to support the Technology Annex to the National Defense Strategy.**
 - o The Secretary of Defense should direct USD (R&E) to produce unclassified emerging technology R&D objectives and publish these objectives publicly. The objectives should represent an unclassified component of the Technology Annex, and be regularly updated as living documents.
 - The R&D objectives should be tied to subsets or components of priority technologies and applications on which the government envisions the private sector playing a major role in building future capabilities.¹⁴⁴ They should be communicated with an appropriate level of detail to provide current defense companies guidance to steer their internal R&D investments, communicate to startups interested in working with the government where future opportunities lie, and signal to venture capitalists where future DoD funding might flow.
 - USD (R&E) should incorporate these objectives into the go-to-market strategy, coordinating digital innovation initiatives to act as surrogates to amplify this communication, and where appropriate, execute these priorities.

- o The Secretary of Defense should direct the Steering Committee on Emerging Technology to develop an appropriate approach to monitor industry-independent R&D investments to gauge the effectiveness of these efforts. This should be coordinated with the DoD Office of General Counsel and relevant industry associations.
- o OUSD (R&E) should leverage public-private exchange programs,¹⁴⁵ as well as internal technical expertise from entities like DARPA and the interagency technology scouting community, to bring both technical expertise and commercial proficiency to the effort.¹⁴⁶

Blueprint for Action: Chapter 2 - Endnotes

¹ See Chapter 9 of this report and its associated Blueprint for Action for a recommendation to establish a Technology Competitiveness Council that could serve as a body for this kind of strategic-level coordination.

² This action is mirrored in the Chapter 3 and Chapter 5 Blueprints for Action. The Commission acknowledges section 236 of the FY2021 National Defense Authorization Act, which permits the Secretary of Defense to establish a steering committee on emerging technology and national security threats composed of the Deputy Secretary of Defense; the Vice Chairman of the Joint Chiefs of Staff; the Under Secretary of Defense for Intelligence and Security; the Under Secretary of Defense for Research and Engineering; the Under Secretary of Defense for Personnel and Readiness; the Under Secretary of Defense for Acquisition and Sustainment; the Chief Information Officer; and such other officials of the Department of Defense as the Secretary determines appropriate. However, the structure described in section 236 does not include leadership from the Intelligence Community, and will thus not drive the intended action. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021), <https://docs.house.gov/billsthisweek/20201207/CRPT-116hrpt617.pdf>.

³ The Commission's recommendation for the development of a Technology Annex to the National Defense Strategy is discussed later in this blueprint.

⁴ For example, the DoD's Joint AI Center (JAIC) is building a Joint Common Foundation (JCF) that aims to provide policies and tools that support an enterprise cloud-enabled AI environment. See *About the JAIC*, JAIC (last accessed Feb. 2, 2021), <https://www.ai.mil/about.html>. Other digital ecosystem pathfinders include, but are not limited to, Platform One, Kessel Run, Space CAMP, Black Pearl, CReATE, ADVANA, and the Army Futures Command Software Factory.

⁵ See DoD Digital Modernization Strategy, U.S. Department of Defense at 30, 42-43 (July 12, 2019), <https://media.defense.gov/2019/Jul/12/2002156622/-1/-1/1/DOD-DIGITAL-MODERNIZATION-STRATEGY-2019.PDF> (describing how the DoD plans to deploy an end-to-end identity, credential, and access management infrastructure). This is an essential function that must be implemented in an interoperable way across the national security-wide digital AI R&D ecosystem. DoD plans include a goal to "Improve and Enable Authentication to DoD Networks and Resources through Common Standards, Shared Services, and Federation." *Id.* at 30.

⁶ Implemented as applications as a service (AaaS).

⁷ Implemented as platforms as a service (PaaS).

⁸ The digital ecosystem should incorporate DevSecOps processes and tools laid out in the DoD Enterprise DevSecOps Reference Design. See DoD Enterprise DevSecOps Reference Design, U.S. Department of Defense (Aug. 12, 2019), https://dodcio.defense.gov/Portals/0/Documents/DoD%20Enterprise%20DevSecOps%20Reference%20Design%20v1.0_Public%20Release.pdf. For more information, see *Understanding the Differences Between Agile & DevSecOps - from a Business Perspective*, GSA (last accessed Jan. 1, 2021), <https://tech.gsa.gov/guides/understanding-differences-agile-devsecops/> ("DevSecOps improves the lead time and frequency of delivery outcomes through enhanced engineering practices; promoting a more cohesive collaboration between Development, Security, and Operations teams as they work towards continuous integration and delivery.").

⁹ For a short primer on MLOps, see *2021 Technology Spotlight - The Emergence of MLOps*, Booz Allen Hamilton (2021), https://www.boozallen.com/content/dam/boozallen_site/dig/pdf/white_paper/the-emergence-of-mlops.pdf.

¹⁰ Notably, the Office of the Under Secretary of Defense for Research & Engineering (OUSD (R&E)) in 2020 outlined a similar vision for an enterprise-wide, shared digital ecosystem to implement the Department's Digital Engineering Strategy and accelerate broad adoption of model-based system engineering. See Andrew Monje, *Future Direction of Model-Based Engineering Across the Department of Defense*, U.S. Department of Defense (Jan. 27, 2020), <https://ac.cto.mil/wp-content/uploads/2020/05/RAMS-Monje-27Jan2020-Future.pdf>.

¹¹ A common software delivery platform used by industry and academia based on the features of Git (<https://git-scm.com>), GitHub (<https://github.com>), and GitLab (<https://about.gitlab.com>).

¹² Implemented as software as a service (SaaS).

¹³ See Chapter 7 of this report. See also Tab 1 - Recommendation 6: "Expedite the development of tools to create tailored AI test beds supported by both virtual and blended environments" in *Second Quarter Recommendations*, NSCAI at 14 (July 2020), <https://www.nscai.gov/previous-reports/>.

¹⁴ See Tab 1 - Recommendation 1: "Create an AI software repository to support AI R&D" in *Second Quarter Recommendations*, NSCAI at 3 (July 2020), <https://www.nscai.gov/previous-reports/>; see also Tab 1 - Recommendation 2: "Promote ATO reciprocity as the default practice within and among programs, Services, and other DoD agencies to enable sharing of software platforms, components, infrastructure, and data for rapid deployment of new capabilities" in *Second Quarter Recommendations*, NSCAI at 5 (July 2020), <https://www.nscai.gov/previous-reports/>.

¹⁵ Implemented as data as a service (DaaS). See Tab 1 - Recommendation 3: "Create a DoD-wide AI data catalog to enable data discovery for AI R&D" in *Second Quarter Recommendations*, NSCAI at 7 (July 2020), <https://www.nscai.gov/previous-reports/>.

¹⁶ The data services and resources provided by the digital ecosystem should support the DoD Data Strategy. See *Executive Summary: DoD Data Strategy*, U.S. Department of Defense (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

¹⁷ For more information on AI documentation, see Chapter 7 of this report and the Appendix containing the abridged version of NSCAI's Key Considerations for Responsible Development & Fielding of AI. See also the Commission's recommendation to produce documentation of the AI life cycle in the section on "Engineering Practices" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁸ As part of the funding plan, the Department should consider proposing expansion of the pilot for consumption-based solutions outlined in Section 834 of the FY2021 NDAA to extend across the stack of managed services that compose the digital ecosystem. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁹ The DoD CIO should determine the appropriate structure for such a working group, and may decide to leverage or federate existing cross-functional working groups such as those for the DoD Enterprise DevSecOps Initiative or Enterprise Infrastructure. Similarly, DoD CIO should work with pathfinder organizations to determine whether they should be incorporated as part of the governance working group or broken out as a separate community from which to draw best practices.

²⁰ For example, contributions to the digital ecosystem would come from AI delivery teams at the combatant command headquarters level, and from forward-deployed teams, as they leverage the ecosystem for agile development of AI-driven capabilities.

²¹ The list included is intended as a general outline of key stakeholders; it is not exhaustive.

²² In recent years, the Department has made promising initial steps to establish managed services constructs for platforms, cloud infrastructure, and software development; for example, the Air Force's CloudOne and Platform One as well as multiple in-house software factories such as Kessel Run and Space CAMP (<https://software.af.mil/software-factories> and <https://software.af.mil/dsop/services/>); the Navy's Black Pearl (<https://blackpearl.us/>); and the Army's Coding Repository and Transformation Environment (CReATE); and the new Software Factory at Army Futures Command. Further, the Office of the Secretary of Defense has built a data management platform, ADVANA, with the goal to establish it as the single authoritative source for audit and business data analytics. See Written Statement for the Record of David L. Norquist, Deputy Secretary of Defense before the U.S. Senate Armed Services Committee Subcommittee on Readiness at 6 (Nov. 20, 2019), https://www.armed-services.senate.gov/imo/media/doc/Norquist_11-20-19.pdf.

²³ The digital ecosystem's open architecture should be developed with consideration of existing warfighting architectures; for example, the Joint Warfighting Network Architecture recommended in Chapter 3 of this report.

²⁴ The open architecture should: 1) Define a common set of well-documented common interfaces for the ecosystem's key components and building blocks; 2) Support and integrate the work of existing pathfinders up and down the ecosystem technology stack; and 3) Incorporate the process elements of the DoD DevSecOps Reference Design Version 1.0 Aug. 12, 2019, data authorizations, and continuous software ATO reciprocity.

Blueprint for Action: Chapter 2 - Endnotes

²⁵ The CIO should include guidance along with the open architecture describing what categories of systems are to be adherent and which may be exempt.

²⁶ Later in this blueprint, NSCAI recommends that USD (R&E) be appointed co-chair and chief science advisor to the Joint Requirements Oversight Council (JROC) for joint and cross-domain capabilities. This recommendation is also emphasized in Chapter 3 of this report.

²⁷ The Executive Summary that accompanies the DoD Data Strategy states that each Component will develop “measurable Data Strategy Implementation Plans, overseen by the CDO and DoD Data Council.” See *Executive Summary: DoD Data Strategy*, U.S. Department of Defense (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

²⁸ For example, ensuring contract Data Item Descriptions include the use of application programming interfaces as the data transfer medium. For additional details on AI documentation, see Chapter 7 of this report and the Appendix containing the abridged version of NSCAI’s Key Considerations for Responsible Development & Fielding of AI. See also the Commission’s recommendation to produce documentation of the AI life cycle in the section on “Engineering Practices” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

²⁹ Memorandum from Deputy Secretary of Defense, *Actions to Enhance and Accelerate Enterprise Data Management*, U.S. Department of Defense at 1 (Dec. 10, 2020).

³⁰ As an example, the Steering Committee on Emerging Technology could consider designating the Defense Logistics Agency (DLA) as an enterprise service provider for logistics applications and associated services. These applications would be made available within the ecosystem for reuse and integration. Similarly, upon publication of the reference design, the Committee could explore working with the Intelligence Community to designate and integrate Intelligence Community application service providers (e.g., the National Geospatial Agency for GEOINT application services).

³¹ The repository would support implementation of the DoD Data Strategy; the Strategy’s guiding principles include “data is a strategic asset” and “enterprise-wide data access and availability.” See *DoD Data Strategy*, U.S. Department of Defense at 3-4 (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

³² See “*Advana*” *Defense Analytics Platform – Department of Defense*, ACT-IAC (June 16, 2020), <https://www.youtube.com/watch?v=BIQ31B9Hv44>.

³³ The digital ecosystem rides on the capacity of DoD’s underlying network and communication backbone to provide rapid, on-demand access to the essential AI building blocks. The DoD Digital Modernization Strategy calls out the need to modernize the Department’s primary networks, the warfighter’s communication connectivity, and coalition networks—highlighting the need to upgrade the optical network transport, routers, switches, and satellite gateways. See *DoD Digital Modernization Strategy*, U.S. Department of Defense at 20-21, 35, 37 (July 12, 2019), <https://media.defense.gov/2019/Jul/12/2002156622/-1/-1/1/DOD-DIGITAL-MODERNIZATION-STRATEGY-2019.PDF>.

³⁴ Components of which are already underway as a result of the JAIC’s Joint Common Foundation initiative.

³⁵ A common software delivery platform used by industry and academia based on the features of Git (<https://git-scm.com>), GitHub (<https://github.com>), and GitLab (<https://about.gitlab.com>).

³⁶ See Tab 1 - Recommendation 1: “Create an AI software repository to support AI R&D” in *Second Quarter Recommendations*, NSCAI at 3 (July 2020), <https://www.nscai.gov/previous-reports/>.

³⁷ See Chapter 7 of this report. See also Tab 1 - Recommendation 6: “Expedite the development of tools to create tailored AI test beds supported by both virtual and blended environments” in *Second Quarter Recommendations*, NSCAI at 14 (July 2020), <https://www.nscai.gov/previous-reports/>.

³⁸ A federated repository system uses a federated directory that ties distributed repositories together as a virtual whole. See Tab 1 - Recommendation 3: “Create a DoD-wide AI data catalog to enable data discoverability for AI R&D” in *Second Quarter Recommendations*, NSCAI at 7 (July 2020), <https://www.nscai.gov/previous-reports/>.

³⁹ This would be supported by the prototype centralized data repository recommended above, and hinges on implementation of the new DoD Data Strategy, which details the goals to make DoD data visible, accessible, understandable, linked, trustworthy, interoperable, and secure. *DoD Data Strategy*, U.S. Department of Defense at 6 (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

⁴⁰ Trained AI models are a special class of data, and the same federated repository system solution used for DataEx can also be used to support ModelEx.

⁴¹ Another type of anticipated trained AI model is digital twins, as modeling and simulation platforms, such as the Army's One World Terrain advance to support training digital twins through reinforcement learning. For more on *One World Terrain*, see *One World Terrain: A Pillar of the Army's Synthetic Training Environment*, USCICT (Aug. 2, 2019), <https://www.youtube.com/watch?v=K50eL1wU4ic>.

⁴² DoD users could choose services off a multi-cloud provider schedule paying as they go for computing resources and uploading/storing the government's data. To facilitate seamless migration of data and software from one platform to another, the DoD should negotiate contracts with providers that appropriately limit expenses related to data egress and migration.

⁴³ Internally developed software solutions and data sets could be contributed for use across the DoD with built-in incentives for contributors through awarded cloud credits when products are contributed and used. Limited public-facing elements could be brokered on the National AI Research Resource recommended in Chapter 11 of this report.

⁴⁴ See Chapter 7 of this report.

⁴⁵ These platform environments should have ATO reciprocity for the building blocks they provision, including incorporating DevSecOps development stacks.

⁴⁶ Digital ecosystem contributors and users include embedded development teams working at the tactical edge (see below Recommendation: Embed AI development capabilities in support of operations); private sector partners contributing trained models; academic researchers working on open challenge problems; researchers working within a DoD lab; or international partners co-developing interoperable AI capabilities.

⁴⁷ Becky Frankiewicz & Tomas Chamorro-Premuzic, *Digital Transformation Is About Talent, Not Technology*, Harvard Business Review (May 6, 2020), <https://hbr.org/2020/05/digital-transformation-is-about-talent-not-technology>.

⁴⁸ As recommended in Chapter 7 of this report, national security departments and agencies should provide ongoing training to help the workforce better interact, collaborate with, and be supported by AI systems—including understanding AI tools.

⁴⁹ Steve Blank & Pete Newell, *What Your Innovation Process Should Look Like*, Harvard Business Review (Sept. 11, 2017), <https://hbr.org/2017/09/what-your-innovation-process-should-look-like>.

⁵⁰ Andrew Moore, et al., *The AI Stack: A Blueprint for Developing and Deploying Artificial Intelligence*, Proc. SPIE 10635 (May 4, 2018), <https://doi.org/10.1117/12.2309483>. For a graphical depiction of the AI stack, see *About*, Carnegie Mellon University Artificial Intelligence (last accessed Jan. 1, 2021), <https://ai.cs.cmu.edu/about>.

⁵¹ Pub. L. 109-364, John Warner National Defense Authorization Act for Fiscal Year 2007, 109th Congress (2006).

⁵² *DoD Instruction 1300.19: DOD Joint Officer Management Program*, U.S. Department of Defense at 14 (April 3, 2018), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/130019p.pdf?ver=2018-04-03-114842-923>.

⁵³ See Chapter 2 of this report and this associated Blueprint for Action section below about leveraging public-private talent exchanges to infuse technical expertise into the acquisition corps for NSCAI's recommendation to create a technology fellows program to support development of a Technology Annex to the National Defense Strategy; there are numerous extant fellowships across the DoD involving emerging technologies.

Blueprint for Action: Chapter 2 - Endnotes

⁵⁴ This action aligns with the recommendation to establish a strategic data node within the digital ecosystem discussed earlier in this Blueprint and with the DoD Data Strategy, which lists Senior Leader Decision Support and Business Analytics as initial areas of focus. See *DoD Data Strategy*, U.S. Department of Defense at 11 (Sept. 30, 2020), <https://media.defense.gov/2020/Oct/08/2002514180/-1/-1/0/DOD-DATA-STRATEGY.PDF>.

⁵⁵ Notably, this recommendation is aligned with Section 836 of the FY 2021 National Defense Authorization Act, which directs the Secretary of Defense to develop and integrate advanced digital data management and analytics capabilities that integrate all aspects of the defense acquisition system; facilitate the management and analysis of all relevant data; enable the use of such data to inform further development, acquisition, management, and oversight of such systems, including portfolio management; and include software capabilities to collect, transport, organize, manage, make available, and analyze relevant data throughout the life cycle of defense acquisition programs. The section further requires capability demonstrations and revised policies to promote the use of digital management and analytics capabilities by March 15, 2022. Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁵⁶ For example, critical human resource variables such as performance and retention are likely related to budget variables (pay), health care variables (accessibility and quality of care) and even real estate variables (housing). These relationships will become transparent and quantifiable when data tagging supports cross domain analyses.

⁵⁷ Gamechanger is an AI-enabled tool designed to analyze documentation of U.S. government legislation, policies, and regulations for semantic content, to trace authorities and responsibilities across documents, and to map authorities and responsibilities to agencies and officials designated therein.

⁵⁸ See discussion below for details on the responsibilities envisioned for JAIC as the Department's AI Accelerator.

⁵⁹ For example, DIU is currently pursuing a number of AI projects to optimize business processes in the DoD—ranging from using AI-driven Robotic Process Automation to reduce labor costs for the Army Comptroller, to improving Air Force readiness with AI-driven predictive maintenance, to leveraging AI-constructed knowledge graphs to rapidly identify supply chain risks. See *JAIC Partners with DIU on AI/ML Models to Resolve Complex Financial Errors*, JAIC (Oct. 1, 2020), https://www.ai.mil/blog/10_01_20-jaic-partners-with-diu-on-ai/ml-models-to-resolve-complex-financial-errors.html; *U.S. Defense Department Awards C3.ai \$95M Contract Vehicle to Improve Aircraft Readiness Using AI*, Business Wire (Jan. 15, 2020), <https://www.businesswire.com/news/home/20200115005413/en/US-Defense-Department-Awards-C3.ai-95M-Contract-Vehicle-to-Improve-Aircraft-Readiness-Using-AI>; *Accrete.AI Accelerates Growth and Product Adoption with Defense Innovation Unit Contract*, Accrete.ai (April 23, 2020), <https://blog.accrete.ai/newsroom/accrete.ai-wins-million-dollar-contract-with-the-defense-innovation-unit>.

⁶⁰ This should include an evaluation of existing policies and regulations on contract data rights, data format, data definitions, and data environments to accelerate application of commercial AI for acquisition, management, and oversight and maximize insights derived.

⁶¹ For a glimpse into the DoD's innovation ecosystem, see *Tap the Innovation Ecosystem*, MITRE: Acquisition in the Digital Age, (last accessed Jan. 25, 2020), <https://aida.mitre.org/demystifying-dod/innovation-ecosystem/>; *Understanding the DoD Innovation Ecosystem*, MITRE: Bridging Innovation (last accessed Jan. 25, 2020), <https://bridge.mitre.org/visualization/>.

⁶² See *Interim Report*, NSCAI at 31 (November 2019), <https://www.nscai.gov/previous-reports/>.

⁶³ The term "digital innovation initiatives" is used here to describe the various entities across the Office of the Secretary of Defense and the military services, such as the Defense Innovation Unit (DIU), AFWERX, NavalX, and Army Applications Laboratory (AAL), that are focused on bridging the gap with the commercial technology section—especially startups and non-traditional vendors—and accelerating the delivery of best-of-breed technology solutions.

⁶⁴ As the Department's Chief Technology Officer, USD (R&E) has both the authority and mandate to coordinate discrete efforts across OSD and the military services to accelerate the adoption of digital technology and expand the national security innovation base (NSIB). However, USD (R&E) must ensure close coordination with USD (A&S) and, in the case of IT and information systems, DoD CIO, to improve the transition of solutions emerging from these organizations into operational systems.

⁶⁵ For example, through current SBIR “bridging” funds described in Chapter 11 of this report or technology-specific supplemental funding recommended later in this Blueprint for Action under, “Make supplemental funding available to drive operational prototyping, scale, and transition of AI technologies.”

⁶⁶ As the Chief Technology Officer of the DoD, USD (R&E) has a “mission to advance technology and innovation.” Additionally, USD (R&E) is responsible for advis[ing] the Secretary of Defense on all matters related to research; engineering; manufacturing; developmental test & evaluation; and technology development, innovation, and protection activities and programs in the DoD and occurring internationally [as well as] establishing priorities across those matters to ensure conformance with Secretary of Defense policy and guidance. For a full list of USD (R&E)’s responsibilities and functions, see *DoDD Directive 5137.02: Under Secretary Of Defense For Research And Engineering (USD (R&E))*, U.S. Department of Defense (Jan. 4, 2021), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/513702p.pdf?ver=2020-07-15-124712-047>. [hereinafter DoDD 5137.02]

⁶⁷ For example, DIU leverages Other Transaction Authority (OTA) and the Commercial Solutions Opening process to “test, field, and scale commercial technology in less than 24 months.” The Air Force’s AFWERX, in partnership with Air Force Research Lab (AFRL) and DIU’s National Security Innovation Network (NSIN), has pioneered new approaches to Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) funding to “increase the efficiency, effectiveness, and transition rate” of the program. See *Annual Report 2019*, Defense Innovation Unit at 4 (2019), https://assets.ctfassets.net/3nanhbtkr0pc/ZF9fhsMe6jtX15APMLall/cd088a59b91857c5146676e879a615bd/DIU_2019_Annual_Report.pdf; *SBIR Open Topics*, U.S. Air Force AFWERX (last accessed Jan. 1, 2021), <https://www.afwerx.af.mil/sbir.html>.

⁶⁸ Also informed by the JAIC, and made accessible through the digital ecosystem.

⁶⁹ As described in Chapter 3 of this report, there should be a push-pull relationship between innovative technologies and concepts such that the Technology Annex informs, but does not limit, the scope of activity. Digital Innovation Initiatives will likely continue to have responsibilities outside of this go-to-market strategy; for example, the acceleration of commercial AI applications for core business processes.

⁷⁰ The impact and potential use cases of investments may not be apparent for several years. This review aims to provide insight into current activities so as to influence, but not dictate, modifications to the next “go-to-market strategy.” This process should be automated to the maximum extent possible to minimize overhead.

⁷¹ Many of the processes and technical roadblocks faced by traditional and non-traditional vendors that slow critical efforts to build and integrate AI systems will be greatly diminished by the implementation of a digital ecosystem, as described above. However, until then, top-down support at the highest levels of leadership will be essential to empower digital innovation initiatives. Per DoDD 5137.02, part of USD (R&E)’s functions include working in conjunction with the USD (A&S) to identify, evaluate, and promote opportunities to reduce barriers to entry for commercial technologies and non-traditional defense partners; and leading initiatives to engage non-traditional suppliers of technology. See DoDD 5137.02.

⁷² Where appropriate, efforts should leverage expertise from FFRDCs and UARCs.

⁷³ Prototyping contracts provide non-recurring engineering dollars to companies for early-stage technologies and projects “to evaluate and inform [their] feasibility or usefulness.” Often, these dollars come from dedicated funds, such as the SBIR and STTR programs and DIU’s internal prototyping budget; and are distributed by organizations like DIU outside of the acquisition life cycle domains, including DoD programs of record (PoR). Companies executing promising projects through these mechanisms often exhaust prototype funding and are unable to secure sustainable follow-on contracts (i.e., with a clear path toward integration into a PoR) because they cannot identify a customer, or the customer’s funding is already committed. AFWERX improved transition in its SBIR program by achieving buy-in from potential customers through matching program funds. See Tab 1 - Recommendation 7: “Strengthen Return on SBIR Investments” in *Interim Report and Third Quarter Recommendations*, NSCAI at 52 (October 2020), <https://www.nscai.gov/previous-reports/>; Prototyping Guidebook, U.S. Department of Defense at 36 (November 2019), <https://www.dau.edu/tools/Lists/DAUTools/Attachments/329/DoD%20Prototyping%20Guidebook.%20v2.0.pdf>.

Blueprint for Action: Chapter 2 - Endnotes

⁷⁴ For example, at least one F-22 and F-35 aircraft designated as AI test beds could incentivize existing contractors and non-traditional firms to work together and better align their incentives to field new mission capabilities. Such an initiative would build on initial efforts to integrate agile software development into F-22 modernization programs. See Craig Ulsh, *Software Acquisition and Practices (SWAP) Study: Vignettes*, DoD Defense Innovation Board at 6 (Jan. 10, 2019), https://media.defense.gov/2019/Mar/07/2002097482/-1/-1/0/SWAP_STUDY_VIGNETTES.PDF.

⁷⁵ The 2019 National Defense Authorization Act identified metrics for DIU to report, such as: the number and types of transitions by the Unit to the military departments or fielded to the warfighter; and the impact of the Unit's initiatives, outreach, and investments on Department of Defense access to technology leaders and technology not otherwise accessible to the Department, including the number of non-traditional defense contractors with Department of Defense contracts or other transactions resulting directly from the Unit's initiatives, investments, or outreach; the number of traditional defense contractors with contracts or other transactions resulting directly from the Unit's initiatives; and the number of innovations delivered into the hands of the warfighter. See Pub. L. 115-232, sec. 244, John S. McCain National Defense Authorization Act for Fiscal Year 2019, 115th Congress (2018); Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁷⁶ This metric should be appropriately scoped such that consideration is given to products or solutions that lend themselves to enterprise licensing agreements and prioritize measures that indicate the level of cross-service, cross-unit proliferation of a solution.

⁷⁷ A 2019 study conducted by the Defense Innovation Board Defense reached similar conclusions with regard to software acquisitions generally, stating that the current approach to software development is broken and is a leading source of risk to DoD; it takes too long, is too expensive, and exposes warfighters to unacceptable risk by delaying their access to tools they need to ensure mission success. *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, DoD Defense Innovation Board at i (May 2019), <https://media.defense.gov/2019/May/01/2002126689/-1/-1/0/SWAP%20COMPLETE%20REPORT.PDF>.

⁷⁸ The Adaptive Acquisition Framework promotes use of tailored acquisition approaches based on the needed capability. It includes six guiding pathways for acquiring capabilities that Milestone Decision Authorities (MDAs), other Decision Authorities (DAs), and Program Managers (PMs) can "tailor, combine, and transition between": Urgent Capability Acquisition, Middle Tier of Acquisition, Major Capability Acquisition, Software Acquisition, Defense Business Systems, and Acquisition of Services. See *Adaptive Acquisition Framework Pathways*, Defense Acquisition University, (last accessed Dec. 26, 2020), <https://aaf.dau.edu/aaf/aaf-pathways/>. The Software Acquisition Pathway was developed based on a recommendation from the Defense Innovation Board in the 2019 Software Study. See *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, DoD Defense Innovation Board at 37, S2 (May 2019), <https://media.defense.gov/2019/May/01/2002126689/-1/-1/0/SWAP%20COMPLETE%20REPORT.PDF>.

⁷⁹ The Contracting Cone outlines all Federal Acquisition Regulation (FAR) and Non-FAR contract strategies. *Contracting Cone*, Defense Acquisition University (last accessed Dec. 20, 2020), <https://aaf.dau.edu/aaf/contracting-cone/>.

⁸⁰ Specialized pathways include approaches captured within the Department's Adaptive Acquisition Framework such as the Middle Tier of Acquisition and Software Acquisition that are exempted from certain requirements within the Defense Acquisition System.

⁸¹ In January 2020, former Under Secretary of Defense for Policy Michele Flournoy cited concern over inadequate training and incentives for acquisition professionals to make full use of authorities provided by Congress. She noted "pockets of [acquisition] excellence" in Special Operations Command and the Air Force, but argued the larger acquisition corps "is not using the authorities effectively, consistently, and at scale." See Testimony of The Honorable Michele A. Flournoy, former Under Secretary of Defense for Policy before the U.S. House of Representatives Armed Services Committee, *Hearing on DoD's Role in Competing with China* at 6 (Jan. 15, 2020), https://armedservices.house.gov/_cache/files/4/4/44fbef3d-138c-4a0a-b3a9-2f05c898578f/0E4943A5BFAE DA465D485A166FABCF5F.20200115-hasc-michele-flournoy-statement-vfinal.pdf.

⁸² Including Federal Acquisition Regulation (FAR)-based approaches and non-FAR-based approaches as outlined in the Defense Acquisition University's Contracting Cone. See *Contracting Cone*, Defense Acquisition University (last accessed Dec. 20, 2020), <https://aaf.dau.edu/aaf/contracting-cone/>.

⁸³ Such as the middle tier of acquisition and the software acquisition pathway.

⁸⁴ For example, efforts associated with section 230 of the Fiscal Year 2020 NDAA on talent management of digital expertise and software professionals; section 256 on an education strategy for Artificial Intelligence; and section 862 of the FY2020 NDAA on software development and software acquisition training and management programs. In support of the implementation of Section 862, USD (A&S) is developing a pilot software acquisition training program that aims to better enable the “creation and execution of acquisition strategies and contracts that support the speed of technology and change” by providing students with the foundations of digital technologies through evolutionary content in context of the Defense Acquisition System. *Digital DNA: Software Acquisition Training Pilot*, U.S. Department of Defense at 1 (on file with the Commission); see also *Report to Congress on FY20 NDAA Section 862(b)(1)(B) Software Development and Software Acquisition Training and Management Programs*, U.S. Department of Defense at Appendix H (January 2021), https://www.hci.mil/docs/Policy/FY20_NDAASec862ReportToCongress_DoDSoftwDevSoftwAcqTngMgt_Jan2021.pdf.

⁸⁵ This should be coordinated appropriately with the relevant legal and ethics officials to avoid any potential conflicts of interest.

⁸⁶ Section 1102 of the FY2021 National Defense Authorization Act directs the Secretary of Defense to provide briefings to the defense authorization committees on implementation of public-private exchange programs and recommendations for statutory changes to improve their use and effectiveness. Section 1102 also directs the Secretary to take steps to ensure the exchange program is applied to the defense modernization priorities—including AI. While USD (R&E)’s modernization directors are responsible for “unifying and advancing the Department’s investments and capabilities [in their areas], and ensur[ing] the transition of technologies into operational use,” the Department’s acquisition professionals will be the personnel ultimately responsible for operationalizing the modernization priorities. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021); see also *Modernization Priorities*, U.S. Department of Defense, USD (R&E), (last accessed Dec. 28, 2020), <https://www.cto.mil/modernization-priorities/>.

⁸⁷ The DoD has already begun to make progress in these areas. For example, the Advanced Distributed Learning (ADL) Initiative under the Office of the Under Secretary of Defense for Personnel and Readiness (USD (P&R)), which originated in the 1990s, is a DoD-wide program for modernizing DoD training/education, including the use of learning technologies and platforms, and support for content sharing, collaboration, and interoperability. ADL is currently pursuing an Enterprise Course Catalog to federate disparate or decentralized catalogs across the organization, aggregating the content into a single, Defense-wide portal. See *Enterprise Course Catalog (ECC)*, Advanced Distributed Learning Initiative (last accessed Feb. 12, 2021), <https://adlnet.gov/projects/ecc/>.

⁸⁸ Including DoD-specific training as well as relevant commercial and open-source training.

⁸⁹ Examples could include draft acquisition strategy documents for programs planning to use the middle tier or software acquisition pathways; model contracting language for AI technologies, etc.

⁹⁰ Including on new or innovative acquisition approaches and best practices as well as new or emerging digital technologies and technical approaches (e.g., digital engineering, MLOps, etc.).

⁹¹ This recommendation echoes a recommendation made by the Defense Innovation Board (DIB) in a 2019 study on software acquisition and practices within the Department of Defense. The DIB called for a new acquisition pathway for software that would prioritize continuous integration and delivery of working software in a secure manner, with continuous oversight from automated analytics. The DIB provided draft legislative language in the body of the report for consideration by the DoD and Armed Services Committees in implementing this recommendation. The draft legislative text indicated the need for a rapid contracting mechanism to be established as part of the software pathway. Although the creation of a software acquisition pathway was directed by section 800 of the FY2020 NDAA and the Department has since issued a formal policy on the pathway, the rapid contracting mechanism remains unimplemented. See *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, DoD Defense Innovation Board at S58 (May 2019), <https://media.defense.gov/2019/May/01/2002126689/-1/-1/0/SWAP%20COMPLETE%20REPORT.PDF>; Pub. L. 116-92, National Defense Authorization Act for Fiscal Year 2020; *DoD Instruction 5000.87: Operation of the Software Acquisition Pathway*, U.S. Department of Defense (Oct. 2, 2020), https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/500087p.PDF?ver=virAfQj4v_LgN1JxpB_dpA%3D%3D.

Blueprint for Action: Chapter 2 - Endnotes

⁹² Using system-level program elements incentivizes programs to build “full stack” with each subcomponent and enabling technology being built and procured individually as part of the broader program makeup. This reduces risk. In recent years, enabling services such as Platform One have re-emerged, but it is difficult to justify base operating budgets for these organizations because they are not tied to discrete outputs. See Eric Lofgren, *The DoD Budget Process: the Next Frontier of Acquisition Reform*, George Mason University Center for Government Contracting (July 9, 2020), https://business.gmu.edu/images/GovCon/White_Papers/The_DoD_Budget_Process.pdf.

⁹³ Joint Capabilities Integration and Development System (JCIDS) and the Planning, Programming, Budget and Execution (PPBE) process are tightly linked. Military needs drive the development of new programs to deliver capability. Traditionally derived from concepts of operations, these needs are the basis against which the Department evaluates, costs, and ultimately pursues a solution. If the Department determines that a material solution is necessary, the need will be decomposed into requirements that prescribe the design, specification, and function of the system intended to deliver the capability. Once validated, these requirements drive the DoD's budget. *Id.* at 5.

⁹⁴ Commonly known as “colors of money,” DoD funds are appropriated into the following categories, each with its own allowable uses per law: Research, Development, Test & Evaluation (RDT&E) dollars, Procurement dollars, Operations & Maintenance (O&M), and Sustainment dollars.

⁹⁵ The distinction between research and development funds and operating funds disincentivizes the cycle of continuous development and integration necessary to derive value from AI and software-based applications. Within the RDT&E appropriation alone, separate funding for research, development, prototyping, and fielding assumes a slow linear progression from lab to field pre-defined system requirements that allow for little to no user feedback. Once fielded, appropriations law governing the use of O&M funds challenges upgrades to digital systems.

⁹⁶ Congressional testimony from former Under Secretary of Defense for Policy Michele Flournoy highlights the centrality of experimentation to developing new concepts and capabilities at the speed required to outpace our competitors. See Testimony of The Honorable Michele A. Flournoy, former Under Secretary of Defense for Policy before the U.S. House of Representatives Armed Services Committee, *Hearing on DoD's Role in Competing with China* at 8 (Jan. 15, 2020), https://armedservices.house.gov/_cache/files/4/4/44fbef3d-138c-4a0a-b3a9-2f05c898578f/0E4943A5BFAE DA465D485A166FABCF5F.20200115-hasc-michele-flournoy-statement-vfinal.pdf.

⁹⁷ Requirements are developed that drive technological development, and prototyping and experimentation occur as a means to refine requirements and manage risk. This incentivizes integration of incremental technologies into programs of record rather than disruptive or rapidly changing user-centered technologies, such as AI; and limits the ability of program managers to respond to any fast-paced change in technology later in the life of the program. See Pete Modigliani et al., *Modernizing DoD Requirements: Enabling Speed, Agility, and Innovation*, The MITRE Center for Technology and National Security (March 2020), <https://www.mitre.org/sites/default/files/publications/pr-19-03715-2-modernizing-dod-requirements-enabling-speed-agility-and-innovation.pdf>.

⁹⁸ The budget activity 8 (BA 8) pilot seeks to overcome the barrier that DoD spending categories pose to the development and sustainment of digital technologies. The Office of the Under Secretary of Defense for Acquisition and Sustainment and the Office of the Under Secretary of Defense for Comptroller selected nine programs to begin to pilot the BA 8 for FY2021. Defense appropriators approved eight of the nine programs, and BA 8 is being established for each Service and Defense-wide under the Research, Development, Test & Evaluation appropriation and enable two-year funding. See H.R. 133, Consolidated Appropriations Act, 2021, 116th Congress (2020), <https://docs.house.gov/billsthisweek/20201221/BILLS-116RCP68-JES-DIVISION-C.pdf?source=email> (joint explanatory statement at 118).

⁹⁹ Appointing USD (R&E) Co-Chair and Chief Science Advisor to the JROC would help push forward efforts to reform requirements generation and validation. Serving as the system architect for joint and cross-domain solutions, USD (R&E) would advocate for more flexible system design and specifications such as modular open systems architecture and standards, well-documented application programming interfaces (APIs). See Chapter 3 of this report. See also Tab 2 - Recommendation 2: “USD (R&E) should be appointed the Co-Chair and Chief Science Advisor to the Joint Requirements Oversight Council (JROC) for Joint and cross-domain capabilities” in *Interim Report and Third Quarter Recommendations*, NSCAI at 70 (October 2020), <https://www.nscai.gov/previous-reports/>.

¹⁰⁰ Section 809 of the FY2021 NDAA directs the Secretary of Defense and the Director for Extramural Innovation and Research Activities to “conduct an assessment of the processes for developing and approving capability requirements for the acquisition programs of the Department of Defense and each military department” and submit reports to the defense authorization committees. Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021). Section 809 further stipulates that, as part of the assessment, both officials must evaluate the “extent to which portfolio management techniques are used in the process for development capability requirements to coordinate decisions and avoid duplication of capabilities across acquisition programs.” *Id.* The Joint Explanatory Statement accompanying the provision indicates that the Department shall consider the recommendations made in the MITRE Corporation’s *Modernizing the Requirements Process: Enabling Speed, Agility, and Innovation* as part of the directed assessment. Recommendations include the establishment of enterprise-level requirements or “Warfighter Essential Requirements” for capabilities to ensure acquisition programs are closely aligned to warfighter needs, drive systems of systems approaches and reduce redundancies between and among services and domains; and enable budget and requirements trade-offs through a portfolio management approach. The authors also recommend different management approaches for requirements based on the attributes of the system being developed. See Pete Modigliani, et al., *Modernizing the Requirements Process: Enabling Speed, Agility, and Innovation*, MITRE (March 2020), <https://www.mitre.org/sites/default/files/publications/pr-19-03715-2-modernizing-dod-requirements-enabling-speed-agility-and-innovation.pdf>.

¹⁰¹ A formal legislative proposal may not be required. DoD retains discretion in the structure and objectives of annual budget proposals. However, approval from Congress and the Office of Management and Budget is required.

¹⁰² Such as dashboards and digital engineering artifacts.

¹⁰³ USD (R&E) should work closely with the JAIC, the Joint Staff, and the military services to identify specific programs and mission areas ripe for potential application of AI technologies, with particular attention to near-term warfighter needs from the Combatant Commands, and use the fund to accelerate efforts in those areas. Establishment of this fund would need to be accompanied with transfer authority such that USD (R&E) could transfer resources to the relevant entities to conduct these activities.

¹⁰⁴ This is being led by the DoD Office of the Under Secretary of Defense for Comptroller and Office of the Under Secretary of Defense for Acquisition and Sustainment, based on the findings and recommendations of the Defense Innovation Board’s Software Acquisition and Practices Study. *Software is Never Done: Refactoring the Acquisition Code for Competitive Advantage*, DoD Defense Innovation Board (May 2019), https://media.defense.gov/2019/Apr/30/2002124828/-1/-1/0/FINAL_SWAP_REPORT.PDF. Jeff Boleng, Special Assistant for Software Acquisition to the Under Secretary of Defense for Acquisition and Sustainment, publicly stated the goal of the pilot as “simplifying the budget process, increasing the visibility, accountability of the funding.” Billy Mitchell, *DOD has OMB Support for Special Software-only Appropriations Pilots*, FedScoop (Sept. 10, 2019), <https://www.fedscoop.com/dod-omb-support-special-software-appropriations-pilots/>. In public remarks made March 3, 2020, Under Secretary of Defense for Acquisition and Sustainment Ellen Lord underscored the significance of the pilot, asserting, “we will begin to see results almost instantaneously, because the administrative burden of making sure you are charging the right development number, the right production number, the right sustainment number, slows things down.” Jared Serbu, *Pentagon Teeing Up Nine Programs to Test New ‘Color of Money’ for Software Development*, Federal News Network (March 4, 2020), <https://federalnewsnetwork.com/acquisition/2020/03/pentagon-teeing-up-nine-programs-to-test-new-color-of-money-for-software-development/>; *West 2020: 3 March 2020 Morning Keynote with The Honorable Ellen Lord*, WEST Conference (March 3, 2020), <https://www.youtube.com/watch?v=VGLqjyMhtok&list=PLFZb4znlHwx0TcsirmyYD6k5BAYxDRwU0&index=6&t=0s>.

¹⁰⁵ For example: budget activities within the appropriation could be aligned to a DoD Component; program elements or budget lines under the budget activities would align to joint capabilities (e.g., Joint Command and Control) and then further decomposed into projects (i.e., key systems, investments, and supporting activities).

¹⁰⁶ Often, technology that has been in use in the commercial sector for years.

Blueprint for Action: Chapter 2 - Endnotes

¹⁰⁷ Section 232 of the National Defense Authorization Act for FY2021 designates the JAIC as a direct report to the Deputy Secretary of Defense, adds to the JAIC's responsibilities the "acquisition and development of mature artificial intelligence technologies in support of defense missions," and directs the Secretary of Defense to clarify the roles and responsibilities of various DoD Components relative to the "research, development, prototyping, testing, procurement of, requirements for, and operational use of artificial intelligence technologies." See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁰⁸ Best practices could include user-centered approaches such as problem discovery, which could be captured and shared via a modern, queryable knowledge management system; or algorithms or models added to the JAIC's repository within the digital ecosystem.

¹⁰⁹ For example, identity-based user authentication and access controls; definition of common standard interfaces and documentation requirements; and accreditation and ATO reciprocity. See full list above.

¹¹⁰ For example, working with the Office of the Under Secretary of Defense, the Defense Contract Management Agency, Service Acquisition Executives, and other relevant parties responsible for acquisition and procurement activities to develop model contract language that incorporates the standards and practices outlined in NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. These would apply both to systems developed by DoD, as well as those that are acquired (including Commercial off-the-shelf systems or those developed by contractors). See the Appendix of this report containing the abridged version of NSCAI's Key Considerations for Responsible Development & Fielding of AI. For additional details on the Commission's recommendation for governance, see the sections on "Aligning Systems and Uses with American Values and the Rule of Law" and "Accountability and Governance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹¹¹ More details for a TEVV framework can be found in Chapter 7 of this report.

¹¹² AI applications are extremely diverse and thereby necessitate a wide range of testing methods. Establishing common approaches to tailoring appropriate processes and tools to the type of AI application at hand will support the ability of DoD components to embrace and scale AI solutions by shortening the testing cycle and making test results interpretable and comparable across the Department. Given the diversity of use cases, the framework would not embody a one-size-fits-all approach, but rather provide core capabilities and guidance adaptable across application areas. For a full discussion on this framework, and required resourcing, see Chapter 7 of this report.

¹¹³ Depending on the current state of the implementation of the digital ecosystem, this shared access could be accomplished through the federated system of distributed software repositories—whether the JAIC's software repository or one managed by a DoD component that originally developed or licensed the software tool.

¹¹⁴ Including tools for TEVV. This effort should also determine what AI development tools are already available across the Department (e.g., where commercial software licenses already exist) and, leveraging the acquisition authority granted in the FY2021 NDAA, procuring leading-edge AI development tools with licensing terms to support enterprise-wide usage. Reasonable consideration should be given for the maturity of the product/tool and likelihood of enterprise use. Section 808 of the FY 2021 National Defense Authorization Act grants the Director of the Joint Artificial Intelligence Center acquisition authority up to \$75 million out of the funds made available in FY 2021-2015 to enter into new contracts. Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹¹⁵ Such as data scientists, AI and software engineers, infrastructure engineers, product managers, and other key positions.

¹¹⁶ Including full stack development, engineering, testing, integration, etc., for AI applications and systems.

¹¹⁷ To reduce barriers to entry, the Department could also consider pairing the Blanket Ordering Agreement or Blanket Purchasing Agreement with a Broad Agency Announcement or Commercial Solutions Open solicitation procedure.

¹¹⁸ This should include a diverse cross-section of expertise that at a minimum includes engineering (i.e., data science and AI solutions), AI digital ecosystem architecture, AI software experts, product managers; and acquisition, legal, policy experts as well as domain experts.

¹¹⁹ This could also involve JAIC representatives embedded at Combatant Command headquarters where appropriate and feasible.

¹²⁰ Of note, the NSCAI Interim Report Appendix 3: Workforce Model's recommendations are designed to support this model, with AI experts and developers serving at hubs, developers serving in spokes, and deployment specialist training helping domain experts maintain data sets and software and better partner with experts and developers. *Interim Report*, NSCAI at 61 (November 2019), <https://www.nscai.gov/previous-reports/>.

¹²¹ CCMDs have specific operational needs that routinely outpace centralized development approaches. Modern battlefield dynamics require that each commander have the ability to tailor the character of his or her war to out-adapt the adversary.

¹²² To stand-up these teams quickly, the CCMDs could leverage the enterprise contracting vehicles through the JAIC to access a pre-vetted pool of talent with AI engineering, data science, and product management competencies. If local contracting vehicles are used, contract provisions should require that all development efforts are interoperable with and leverage the digital ecosystem.

¹²³ In this way, the AI delivery teams will contribute to a growing resource of shared data and software within the digital ecosystem by consuming ecosystem services, developing and fielding tailored AI capabilities, and integrating them into sustainable projects available for use across the department.

¹²⁴ As an example, both Army Futures Command (AFC) and Army Special Operations Command (USASOC) use a model known as "tactical data teams." This model brings AI/ML expertise forward to the field in the form of three- to six-person teams to build AI solutions for real-time operational problems. Executed by a small business, Striveworks, under contract with AFC and USASOC, they are currently supporting efforts in Central Command and Indo-Pacific Command Areas of Responsibility.

¹²⁵ These are similar interactions with the digital ecosystem as those taken by the delivery teams at Combatant Command HQ, only the forward-deployed development team will be consuming digital ecosystem services locally on their provisioned mobile platform. Collocation of the developers with operators will drive real-time experimentation and shorten application feedback loops.

¹²⁶ DoD lacks reliable budget data to track its investments in AI and other critical technologies; a weakness that should be addressed at the source with AI applications that assist humans in generating program descriptions and other budget artifacts.

¹²⁷ For a full discussion of how AI will change warfare, see Chapter 3 of this report.

¹²⁸ For a list of priority AI R&D research areas, see Chapter 3 of this report.

¹²⁹ The Defense Science Board has recommended the level of 3.4% to mirror best practices in the private sector multiple times. *Department of Defense Research, Development, Test, and Evaluation (RD&E): Appropriations Structure*, Congressional Research Service at 12 (Oct. 7, 2020), <https://fas.org/sgp/crs/natsec/R44711.pdf>.

¹³⁰ While defense budgets are projected to flatten or decline in the coming years, the threat environment will only increase in complexity. To meet these new realities, we must create more room in the budget while simultaneously increasing the lethality of our forces. By retiring legacy systems and investing more in emerging technologies and, over the longer term, portfolios of attritable systems, DoD can pursue these needs in tandem, boosting the composability and adaptability of our military forces.

Blueprint for Action: Chapter 2 - Endnotes

¹³¹ Former Secretary of Defense Mark Esper pioneered his “night-court” budgeting process as Army Secretary (2017-2019) and later applied it Department-wide. He “took a hard look at legacy department programs and cut a number of them, refocusing funds on efforts to challenge China and Russia.” As Army Secretary, he “helped guide those restructurings through Congress, and the process, which found around \$25 billion in savings, has garnered largely positive reviews.” Aaron Mehta & Joe Gould, *Night Court Comes to the Pentagon*, Defense News (Aug. 28, 2019), <https://www.defensenews.com/pentagon/2019/08/28/night-court-comes-to-the-pentagon/>. According to the Pentagon’s press release detailing the highlights of the FY2021 budget proposal, the process applied defense-wide generated \$5.7 billion in FY2021 savings, \$0.2 billion in Working Capital Fund efficiencies, and another \$2.1 billion in activities and functions realigned to the Services. Press Release, The Office of the Under Secretary of Defense for Comptroller, *DoD Releases FY 2021 Budget Proposal*, https://comptroller.defense.gov/Portals/45/Documents/defbudget/fy2021/fy2021_Press_Release.pdf.

¹³² This echoes a recommendation made by the Future of Defense Task Force. The Task Force recommended that Congress commission the RAND Corporation (or similar entity) and the Government Accountability Office to study legacy platforms within the Department of Defense and determine their relevance and resiliency to emerging threats over the next 50 years. The Task Force further recommended that upon completion of the studies, “a panel should be convened, comprising Congress, DoD, and representatives from the industrial base to make recommendations on which platforms should be retired, replaced, or recapitalized.” *Future of Defense Task Force Report 2020*, House Armed Services Committee at 8 (Sept. 23, 2020), https://armedservices.house.gov/_cache/files/2/6/26129500-d208-47ba-a9f7-25a8f82828b0/424EB2008281A3C79BA8C7EA71890AE9.future-of-defense-task-force-report.pdf.

¹³³ As noted in the discussion above on building a technical backbone, new programs should also adhere to the digital ecosystem reference design.

¹³⁴ See *Interim Report and Third Quarter Recommendations*, NSCAI at 66 (October 2020), <https://www.nscai.gov/previous-reports/>.

¹³⁵ The purpose and proposed contents of such a Technology Annex are discussed below.

¹³⁶ See recommendations for the IC to increase S&T expertise and intelligence collection in Chapter 5 of this report.

¹³⁷ In its response to the 2017 NDAA provision creating USD (R&E), the DoD specified that the new organization would organize around three major themes. The first was an SIAC that would focus on understanding the enemy’s capabilities and vulnerabilities, conducting analysis on our own U.S. capabilities, tracking technology trends across the globe and assessing potential/emerging threats and/or future opportunities that warrant action, that merit investment. However, since the establishment of USD (R&E), the SIAC has been downgraded from a direct report to the Under Secretary and largely focused on examining threat technologies for OSD customers. See *Report to Congress, Restructuring the Department of Defense Acquisition, Technology and Logistics Organization and Chief Management Officer Organization*, U.S. Department of Defense at 8 (August 2017), <https://dod.defense.gov/Portals/1/Documents/pubs/Section-901-FY-2017-NDAA-Report.pdf> [hereinafter 2017 AT&L Reorganization Plan]

¹³⁸ USD (R&E) has the mandate and authority to perform this function. See DoDD 5137.02 at 5-6.

¹³⁹ This is consistent with a recommendation made in Chapter 3 of this report that the DoD should integrate AI-enabled applications into all major Joint and Service exercises and, as appropriate, into other existing exercises, wargames, experiments, and table-top exercises. See also *Second Quarter Recommendations*, NSCAI at 27 (July 2020), <https://www.nscai.gov/previous-reports/>.

¹⁴⁰ This would also directly support objectives of Section 1102 of the FY2021 NDAA with respect to utilization of public-private talent exchanges. Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁴¹ For example, via the go-to-market strategy outlined above.

¹⁴² An unclassified version of the strategy must be communicated externally, to where the bulk of the AI talent resides. Shifting to a more integrated and transparent communication of priorities would enable Defense primes and non-traditionals to plan and invest more to help meet DoD R&D needs. See Tab 1 - Issue 3: "Expanding Industry's Role in DoD's AI R&D to Develop Next-Generation Capabilities" in *Interim Report and Third Quarter Recommendations*, NSCAI at 48 (October 2020), <https://www.nscai.gov/previous-reports/>.

¹⁴³ This could be done via the reference design for the digital ecosystem outlined above. As stated above, adherence to the reference design should be driven top-down via a memorandum from the Secretary of Defense and enforced through the Joint Requirements Oversight Council (JROC).

¹⁴⁴ For example, under microelectronics, this might include advancing AI multi-chip packages, development of quantifiable assurance, 3D chip stacking, photonics, carbon nanotubes, gallium nitride transistors, domain-specific hardware architecture, electronic design automation, and cryogenic computing. As recommended by NSCAI in our *First Quarter Recommendations*. See *First Quarter Recommendations*, NSCAI at 51 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹⁴⁵ This should be coordinated appropriately with the relevant legal and ethics officials to avoid any potential conflicts of interest.

¹⁴⁶ OUSD (R&E) could leverage existing Intergovernmental Personnel Act authorities as well as the pilot Public-Private Talent Exchange Program. See *Department Of Defense Public-Private Talent Exchange (PPTe) Program: Questions/Answers*, DoD Defense Civilian Personnel Advisory Service (Aug. 23, 2018), https://www.dcpas.osd.mil/Content/Documents/PPTeQuestions_Answers23Aug2018.pdf; Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021) (directing that the Department of Defense establish public-private exchange programs to support the defense modernization priorities).

Chapter 3: AI and Warfare

Blueprint for Action

If U.S. forces are not organized, trained, and equipped for a new warfighting paradigm that is emerging because of artificial intelligence (AI) and other emerging technologies, they will be outmatched and paralyzed by the complexity of the future battlefield.

This Blueprint for Action includes five top-line recommendations to achieve military AI readiness and prepare our forces for the future: 1) Drive organizational reforms through top-down leadership; 2) Develop innovative warfighting concepts; 3) Establish AI-readiness performance goals; 4) Develop and fund advanced technologies and R&D; and 5) Promote AI interoperability and the adoption of critical emerging technologies among U.S. allies and partners.

Recommendation

Recommendation: Drive organizational reforms through top-down leadership.

Continuously out-innovating the competition requires strong commitment from the top civilian and military leaders directing the rapid development and adoption of innovative and disruptive approaches to warfare through top-down governance and oversight processes.

Action for the Department of Defense and the Office of the Director of National Intelligence:

- **Establish a Steering Committee on Emerging Technology, tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.¹**
 - o The Secretary of Defense and Director of National Intelligence should issue a directive immediately establishing the senior oversight committee described above.
 - o The Steering Committee on Emerging Technology provides a forum to drive change, focus, and action on emerging technology that otherwise would not be prioritized. It will enhance intelligence analysis related to emerging technology; connect strategic vision to organizational change; focus concept and capability development on emerging threats; guide defense investments that ensure America's strategic advantage against near-peer competitors; and provide the authority to drive technology adoption and application by the Department.
- **Assign the tri-chair Steering Committee on Emerging Technology responsibility for overseeing the development of a Technology Annex to the next National Defense Strategy²**

Actions for the Department of Defense:

- **Ensure all future JAIC Directors are a three-star general or flag officer with significant operational experience who reports directly to the Deputy Secretary of Defense.**³
 - o Three-star leadership allows the JAIC to engage with the services at a senior rank and within their command structure. Operational experience enables the Director to understand how AI can serve operational requirements and better communicate with the services as to how AI meets capability needs.
- **Appoint Under Secretary of Defense for Research and Engineering (USD (R&E)) as the co-chair and chief science advisor to the Joint Requirements Oversight Council.**⁴
 - o To accelerate AI and other emerging technologies for competitive advantage, USD (R&E) must play a central role in connecting technological advancements in research and development to joint operational requirements.

Action for Congress:

- **In the Defense Authorization Act (NDAA) for Fiscal Year 2022, establish a Steering Committee on Emerging Technology and National Security Threats and designate that it be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.**⁵

Recommendation

Recommendation: Develop Innovative AI-enabled Warfighting Concepts, Informed by Experimentation, Wargames and Real-world Exercises.

Battlefield advantage will shift to those who harness superior data, connectivity, compute power, algorithms, and overall system security to new warfighting concepts. Developing new operational concepts requires Services to incentivize experimentation, and foster a culture of “thinking Red”—in other words, considering the strategies of potential adversaries when developing operational concepts.

Actions for the Department of Defense:

- **Develop innovative operational concepts that integrate new warfighting capabilities with emerging technologies.**
 - o The Deputy Secretary of Defense and the Vice Chairman of the Joint Chiefs should issue a memorandum directing Components and Services to develop a complete deterrence concept for systems warfare that leverages human-machine teaming, AI, and associated technology to prevail against intelligent adversary systems of systems.
 - o Under the guidance from the tri-chair Steering Committee on Emerging Technology, USD (R&E) should receive \$5 million for a team (approximately 20 people) in FY2022 funding to research and develop new AI-enabled capabilities for development and testing of advanced operational concepts. This project must be done in conjunction with DARPA and other capability offices to share the costs of filling technological gaps discovered during the analytic process.

- o These operational concepts should be institutionalized in classified DoD documents that drive comprehensive force development and investment prioritization. Confidential demonstrations should be executed to realize the deterrence concept.
- **Integrate AI-enabled applications into all major Joint and Service exercises and, as appropriate, into other existing exercises, wargames, and table-top exercises.**
 - o The Deputy Secretary of Defense and the Vice Chairman of the Joint Chiefs should issue a memorandum calling for inclusion of AI and other emerging technologies into existing exercises, wargames, and table-top exercises. This includes large-scale exercises and smaller, more frequent events at all echelons.
 - o The purpose of this would be to realize connectivity between systems and sensors, rapid data analysis, faster and more informed decision-making, and more distributed operations.
 - o Concept writers should participate in all major technology demonstrations.
 - o Develop performance objectives and associated metrics to assess integration of AI-enabled applications into exercises, wargames, experiments, and TTXs.
- **Incentivize experimentation with AI-enabled applications through the Warfighting Lab Incentive Fund (WLIF).**
 - o DoD should incentivize experimentation with AI applications across the Department at every level possible by establishing either a special category or prioritized evaluation criteria within the WLIF for proposals that incorporate AI applications.
 - Experimentation with AI-enabled applications are particularly well-suited for the space, cyber, and information domains because of the high volumes of 24/7 data generated in these domains.
 - o The Steering Committee on Emerging Technology should provide annual guidance for selection of WLIF proposals for funding based on priorities developed in the Technology Annex to the NDS.
 - o DoD should increase WLIF funding by \$10 million annually specifically for AI-enabled applications.⁶
- **Encourage a culture of “thinking Red.”**
 - o DoD working closely with the Intelligence Community should develop a granular understanding of our main competitors’ approach to systems confrontation. This will help the Department to better understand our competitors’ operational concepts and to eventually avoid battlefield surprise.
 - o Joint Warfare Analysis Center (JWAC) should be the lead to develop competitors’ operational concepts. Estimate \$2.5 million allocation for a 10-week, 10-game series devoted to mastering red thinking.
 - Red-thinking games must: 1) Integrate deterrence-credibility stretch problems from key classified DoD documents; 2) clear denial concepts for our most stressing scenarios; 3) Be conducted with realistic basing and naval posture; 4) the highest standards of incorporating the best available intelligence; 5) the highest standards of AI-enabled modeling and simulation that ingest and mimic red operations; 6) rigorous two-player adjudication with physics-level detail on red capabilities; and 7) rapid turnaround on force development considerations for the Secretary of Defense.

- o The Office of the Secretary of Defense and the Joint Staff should issue a memorandum directing all military educational institutions to foster in their curriculum the culture of “thinking Red.”

Actions for Congress:

- **Congress should appropriate an additional \$17.5 million to DoD’s budget to support innovative concept development.**

Recommendation: Establish AI and digital readiness performance goals.

Recommendation

To drive outcomes and accountability and provide a means for oversight of Department efforts regulated to AI, DoD should establish key performance objectives and accompanying metrics for AI and digital readiness.⁷

Actions for the Department of Defense:

- **By the end of 2021, establish AI and digital readiness performance goals. To achieve more substantial integration of AI across DoD, the Deputy Secretary of Defense should:**
 - o Direct DoD components to assess military AI and digital readiness through existing readiness management forums and processes. The Steering Committee on Emerging Technology should work closely with the Under Secretary of Defense for Personnel and Readiness,⁸ the Joint Staff, and the JAIC to ensure the identified AI and digital readiness performance objectives are incorporated into the military services’ readiness reporting recovery frameworks, and resourcing strategies.
 - o Direct the military services to accelerate review of specific skill gaps in AI, to inform recruitment and talent management strategies and *provide a report within 12 months.*
 - Assess the number of civilian personnel needed in software developer, software engineer, knowledge management, data scientist, and AI career fields for both management and specialist tracks.
 - Assess the number of military personnel needed in software development, data science, and AI career fields, in both management and specialist tracks, and for commissioned and enlisted personnel.
 - Assess the specialties and personnel required for a DoD and military service digital corps.
 - Establish annual retraining and recruiting goals to create and maintain the personnel described above.
- **Direct the military services, in coordination with the Under Secretary of Defense (Acquisition and Sustainment), the Joint Staff, and the Defense Logistics Agency, and enabled by enterprise services and expertise at the JAIC, to prioritize integration of AI into logistics and sustainment systems wherever possible.**
 - o The Deputy Secretary of Defense should issue a memorandum directing the military services to accelerate use of AI and apply commercial best practices in

predictive analytics for maintenance and supply chain to optimize all classes of supply, equipment, and parts.⁹ The Deputy Secretary of Defense should establish a \$100 million fund, administered by the Under Secretary of Defense (Acquisition and Sustainment) to provide matching contributions to service and agency efforts based on estimated financial or operational return on investment.

- o By the end of 2021, the Under Secretary of Defense (Acquisition and Sustainment), supported by Senior Acquisition Executives and in coordination with the DoD CDO and the JAIC, will establish performance objectives and identify best approaches to achieve data-ready systems in logistics and sustainment systems to support application of AI. Disparate conditions of data-readiness in existing and future systems will require differential approaches to achieve AI-readiness. Broadly, these categories of data-readiness are:
 - Systems with proprietary vendor data (ex. F-35 Joint Strike Fighter, M1 Abrams Tank)
 - Systems with government-owned data (ex. Maintenance and Availability Data Warehouse)
 - Systems that are data-ready (government-owned data that has been documented/tagged for storage/discovery and has published schema for data access (ex. Next Generation Air Dominance, T-7 Redtail, Ground Based Strategic Deterrent).

Actions for Congress:

- **Require the Secretary of Defense to establish performance objectives and accompanying metrics for AI and digital readiness, and provide an update to Congress no later than 120 days after approving these goals.**

Recommendation

Recommendation: Develop and Fund Advanced Technologies and R&D.

Development and fielding of advanced AI-enabled technologies will remain a critical component of DoD's ability to achieve decision advantage on the battlefield.

Actions for the Department of Defense:

- **Define a joint warfighting network architecture by the end of 2021.** OSD CIO and the Joint Staff, in coordination with the Services, should issue a memorandum directing the architecture for a secure, warfighting command and control network. A Service-agnostic warfighting network will enable better integration of AI-enabled technologies with current and future weapon systems. The OSD CIO should provide \$5 million to the right entity to accomplish this design.
- **Invest in priority AI R&D with the support areas that could support future military capabilities.** To accelerate adoption of AI in warfighting missions, the Under Secretary of Defense (R&E) should increase investments¹⁰ in the following priority R&D areas to support future AI-enabled warfighting capability. If advanced, this could build near- and long-term AI-driven capabilities for competitive advantage in a future method of conflict defined by AI. These should be viewed as investments in deterrence in the interim—pursuing critical incremental advancements—and in the long term—building new capabilities yet to be determined that will sustain overmatch. Investments should include:

- o USD (R&E), with the support from DARPA, should prioritize AI R&D for the following topics:
 - The future of teaming—to advance human-AI and AI-AI teaming
 - Advanced scene understanding
 - Intelligent edge devices, computing, and networking
 - Robust and resilient AI
 - Testing and Evaluation, Verification and Validation (TEVV)
 - Integrated AI, modeling, and simulation for decision support
 - Autonomous AI systems
 - Toward more general Artificial Intelligence

Recommendation: Promote AI interoperability and the adoption of critical emerging technologies among allies and partners.

Recommendation

America's enduring relationships with allies and partners represent asymmetric advantages over competitors and adversaries. Differential adoption of AI across military alliances and intelligence partnerships creates interoperability risk that threatens allies' political and military cohesion, the resiliency of alliance structures, and the efficacy of coalition operations. The recommendations that follow reflect a holistic approach to furthering cooperation around AI and emerging technologies in the context of defense, intelligence, and security arrangements. They focus on interoperability and improving capacity and capability development to foster competitive military and intelligence advantages.

Component 1: Enhance Five Eyes efforts to achieve interoperable AI systems.

Actions for the Department of Defense and the Office of the Director of National Intelligence:

- **Coordinate with Five Eyes officials to conduct assessments of the comparative strengths and gaps in AI-related technologies and applications among the Five Eyes allies.**
 - o Assessments would evaluate Doctrine, Organization, Training, Materiel, Leadership and Education, Personnel, Facilities and Policy (DOTMLEPF-P) across the alliance for adopting AI, and future plans for AI-enabled warfighting architectures.
- **Coordinate with Five Eyes officials to develop a five-year plan for improving AI interoperability across the Five Eyes alliance.**
 - o Proposed plans should include, among other things, combined research priorities, development objectives, experimentation, methods to facilitate data sharing, use cases, and common standards for TEVV of AI-enabled systems and interoperability standards. It should also include stress tests for supply chains in critical industries and corresponding risk-mitigation measures.

- o In developing plans, Five Eyes leaders should enhance ongoing efforts of the Technical Cooperation Program,¹¹ through the AI Strategic Challenge (AISC), to further align interoperable AI systems.
- o Five Eyes leaders should continue to advance the joint development of intelligence products by expanding efforts to “increase collection access and reliability, improve the quality and quantity of partner data and analysis, align strategic capabilities and emerging technologies, and promote compatibility across digital architectures and analytic tradecraft.”¹²

Actions for the Department of Defense:

- **Direct a series of AI demonstration pilot projects and host an AI wargame and experimentation series.**
 - o Based on the recommended assessments and planning above, the Secretary of Defense should: 1) Direct a series of AI demonstration pilot projects in areas such as predictive maintenance, autonomous logistics, and sensor fusion with Five Eyes partners across the Future Years Defense Program; and 1) host an AI wargame and experimentation series, beginning with Five Eyes allies.

Component 2: Accelerate NATO AI adoption.

NATO and its member states recognize that AI-related technology has transformative potential for collective security. Coordinated, accelerated, responsible adoption of AI must be an urgent priority across the Alliance in order to address the challenge presented by algorithmic warfare.¹³ NATO allies need to dedicate personnel and resources to support the development and operational applications of AI-related, and other Emerging and Disruptive Technologies (EDTs).

Actions for the Departments of Defense and State:

- **Provide clear policy guidance, technical expertise, and resource support to assist and accelerate NATO’s AI-related initiatives to:**
 - o Ensure AI technologies are incorporated into the *NATO Defense Planning Process*, *NATO Warfighting Capstone Concept*, and plans for *Deterrence and Defense of the Euro-Atlantic Area*.
 - o Evaluate DOTMLPF-P for AI adoption and future plans for AI-enabled warfighting architecture and interoperability in allied or coalition environments.
 - o Support and coordinate development and adoption of foundational definitions, operational and data-sharing practices, technical standards, and architectures focused on interoperability, privacy, and responsible, legal deployment of AI.
 - o Ensure the *NATO Science and Technology Strategy* anticipates technological developments to guide NATO and national research and development priorities.
 - o Develop NATO international staff and allied nation technical expertise.
 - o Conduct simulations, wargaming, experimentation, and pilot projects with use cases for data fusion, data exploitation, and interoperability.

- o Assist in the collaboration with partners beyond the NATO Alliance, including industry and academia
- **Develop, with NATO allies, a proposal for an Alliance-wide AI Implementation Strategy deliverable for NATO Heads of State.**
 - o The proposal should build upon key recommendations of the NATO Reflection Group report submitted to the Secretary General,¹⁴ and should provide guidance on the areas identified above.¹⁵

Component 3: Foster the JAIC AI Partnership for Defense (AI PfD) as a critical vehicle to further AI defense and security cooperation.

Launched in 2020, the AI PfD is a DoD-led effort to convene partner nations to “provide values-based global leadership” on adoption of AI in the defense and security context.¹⁶ Current members include Australia, Canada, Denmark, Estonia, Finland, France, Israel, Japan, Norway, South Korea, Sweden, and the United Kingdom.

Action for the Department of Defense:

- **Prioritize and foster the AI PfD as a critical space for democratic allies and partners to work through defense issues on AI.**
 - o The AI PfD can enhance U.S. efforts to accelerate AI adoption across NATO by supporting key foundational efforts related to data governance and management, infrastructure and technical, legal, and ethics expertise. DoD and Congress should provide continued support to enable the AI PfD to further AI cooperation on defense and security with key allies and partners.

Component 4: Incorporate AI into Indo-Pacific defense cooperation efforts.

Increased opportunities exist for collaboration with Quadrilateral Security Dialogue (Quad) partners India, Japan, and Australia, and other nations committed to advancing a free and open Indo-Pacific region.

Actions for the Departments of Defense and State:

- **Build on the Quad framework and negotiate formal AI-related defense and intelligence cooperation agreements in the Indo-Pacific region with Australia, India, and Japan, as well as with New Zealand, South Korea, and Vietnam.**
 - o This could be done in connection with broader conventional defense and intelligence relationships, and existing security cooperation agreements, or in a stand-alone manner, bilaterally or multilaterally. The U.S. Government should also prioritize AI interoperability at ministerial and working level meetings.¹⁷

Component 5: Create an Atlantic-Pacific Security Technology Partnership to improve defense and intelligence interoperability across Europe and the Indo-Pacific.

An Atlantic-Pacific technology partnership would seek to improve capability and interoperability by bringing together technology innovation with allied and partner militaries and intelligence communities, whether in a NATO, coalition, or other multinational context.

Action for the Departments of Defense and State:

- **Advance a deliberate NATO partnership with Indo-Pacific allies and partners for AI-enabled defense cooperation.**
 - o A NATO-Indo-Pacific partnership focused on AI is needed to facilitate early collaboration and lay the groundwork for interoperability among different allied and partner warfighting architectures.
 - o Plans for such a partnership should include guidance from the tri-chair Steering Committee on Emerging Technology for data sharing, common standards, wargame and experimentation, and improving interoperability of AI systems and warfighting architectures.

Component 6: Modify authorities and processes in order to improve DoD's ability to conduct international capability development.

DoD requires more flexibility in its ability to develop, test, and field AI-enabled systems with existing and new foreign partners, both public and private.

Action for Congress:

- **Expand the flexibility and the agility of the Secretary of Defense's authority to engage in cooperative R&D agreements.**
 - o Legislation should permit DoD to pursue cooperative projects with private companies, academic research centers, and defense and non-defense governmental entities within NATO, major non-NATO allies, and friendly foreign countries, without a direct showing to the improvement of conventional defense capabilities.

- o Legislation should also account for partners' non-monetary contributions, including the value of R&D capabilities and the strategic partnerships, when assessing potential projects.

Actions for the Department of Defense:

- **Review and revise policies related to International Armaments Cooperation to provide flexibility for AI and software driven partnerships.**
 - o The review should include policies related to technology transfer, national disclosure, information and equipment use, equitability requirements, funding requirements, and contracting.
 - o DoD should update policies to provide greater delegation of authorities to Military Departments and Defense Agencies to conclude international agreements.
- **Revise DoD Instruction (DoDI) 5530.03, "International Agreements,"¹⁸ to provide appropriate guidance on AI and software-driven partnerships.**
 - o DoDI 5530.03 should be revised to: 1) enable continuous collaboration on evolutionary hardware and software products that need continuous update across research, development, testing, evaluation, and operational deployment with international partners; 2) provide sufficiently flexible entry and exit criteria for all types of international partners (governmental, industry, and academic) to facilitate capabilities, products, knowledge, and services at the point of need; and 3) provide guidance for acceptable thresholds and limits to balance the protection and promotion aspects of AI-related capability development with international partners.

Blueprint for Action: Chapter 3 - Endnotes

¹ The Commission acknowledges section 236 of the FY 2021 National Defense Authorization Act, which permits the Secretary of Defense to establish a steering committee on emerging technology and national security threats composed of the Deputy Secretary of Defense; the Vice Chairman of the Joint Chiefs of Staff; the Under Secretary of Defense for Intelligence and Security; the Under Secretary of Defense for Research and Engineering; the Under Secretary of Defense for Personnel and Readiness; the Under Secretary of Defense for Acquisition and Sustainment; the Chief Information Officer; and such other officials of the Department of Defense as the Secretary determines appropriate. However, the structure described in section 236 does not include leadership from the Intelligence Community and will thus not drive the intended action. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021), <https://docs.house.gov/billsthisweek/20201207/CRPT-116hrpt617.pdf>.

² This action is described in greater detail in the Chapter 2 Blueprint for Action, which designates a member of the Steering Committee on Emerging Technology the Executive Agent responsible for developing the Technology Annex and outlines the recommended contents and use for the Appendix.

³ Notably, section 236 of the FY2021 NDAA designates the Director of the Joint Artificial Intelligence Center as a direct report to the Deputy Secretary of Defense. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁴ This echoes an action in the Chapter 2 Blueprint for Action, which emphasizes that to reduce redundancies, increase interoperability, and drive a system-of-systems approach to requirements development and management, USD (R&E) must have a stronger role in the Joint Requirements Oversight Council.

⁵ As indicated above, DoD and ODNI have the authority to establish such a forum without legislative action. However, codifying it into law will ensure that it is sustained through leadership transitions. The defense committees could consider using the FY2022 NDAA to amend section 236 of the FY2021 NDAA. As written, section 236 only “permits” the establishment of such a committee; additionally, the provision does not clearly denote chairs of the committee and does not include any Intelligence Community representation. This recommendation is also discussed in Chapter 5 of this report. Additionally, Chapters 2 and 5 of this report recommend establishing funds to mature, operationally prototype, and transition exceptionally promising AI-enabled technologies. For DoD, USD (R&E) would control those funds and, for the IC, the ODNI CTO would control those funds. Those investments should be informed by the Steering Committee on Emerging Technology.

⁶ FY2021 O&M funding was \$42.4 million. J7 received 110 proposals for FY2021 WLIF funding and selected 20 experimentation efforts. NSCAI staff discussions with JS/J7.

⁷ General Charles Q. Brown, Jr. & General David H. Berger, *To Compete with China and Russia, the U.S. Military Must Redefine ‘Readiness,’* Washington Post (Feb. 1, 2021), <https://www.washingtonpost.com/opinions/2021/02/01/brown-berger-military-readiness/>.

⁸ “P&R must enable, guide, and assess a strategically ready Department of Defense for employment by the Joint warfighter when and where it is needed, adaptive to the strategic geopolitical and threat environments, and evolving military-technological advances.” *Preserving Our Competitive Advantage, Personnel And Readiness Strategy For 2030*, U.S. Department of Defense at 13 (October 2020), https://prhome.defense.gov/Portals/52/Documents/Strategy/PR_Strategy_FINAL_.pdf?ver=KY6Vacn3kT1Gd9fNxnR34w%3D%3D.

⁹ In the FY2021 NDAA, Title II, section 234, Congress directed “the Secretary of Defense to identify a set of no fewer than five use cases of the application of existing artificial intelligence enabled systems to support improved management of enterprise acquisition, personnel, audit, or financial management functions, or other appropriate management functions, that are consistent with reform efforts that support the National Defense Strategy.” Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹⁰ With additional funding for DoD investments in AI R&D recommended in the Chapter 2 Blueprint for Action.

¹¹ *DoD Instruction 3100.08: The Technical Cooperation Program (TTCP)*, U.S. Department of Defense (Oct. 15, 2018), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/310008p.pdf?ver=2017-11-30-114948-343>.

¹² *The AIM Initiative: A Strategy for Augmenting Intelligence Using Machines*, Office of the Director of National Intelligence at 10 (Jan. 16, 2019), <https://www.dni.gov/files/ODNI/documents/AIM-Strategy.pdf>.

¹³ Memorandum from Robert O. Work, Deputy Secretary of Defense, *Establishment of an Algorithmic Warfare Cross-Functional Team (Project Maven)*, U.S. Department of Defense (April 26, 2017), https://www.govexec.com/media/gbc/docs/pdfs_edit/establishment_of_the_awcft_project_maven.pdf.

¹⁴ *NATO 2030: United for a New Era. Analysis and Recommendations of the Reflection Group Appointed by the NATO Secretary General*, NATO at 29-31 (Nov. 25, 2020), https://www.nato.int/cps/en/natohq/news_179730.htm.

¹⁵ For further detail, see *Interim Report and Third Quarter Recommendations*, NSCAI at 187-195 (October 2020), <https://www.nscai.gov/previous-reports/>.

¹⁶ The AI PfD seeks to align “like-minded nations to promote the responsible use of AI, advance shared interests and best practices on AI ethics implementation, establish frameworks to facilitate cooperation, and coordinate strategic messaging on AI policy.” *Joint Statement*, AI Partnership for Defense (Sept. 15-16, 2020), https://www.ai.mil/docs/AI_PfD_Joint_Statement_09_16_20.pdf. The AI PfD held its second formal dialogue in January 2021. *DoD Joint AI Center Facilitates Second International AI Dialogue for Defense*, JAIC (Jan. 27, 2021), https://www.ai.mil/news_01_27_21-dod-joint_ai_center_facilitates_second_international_ai_dialogue_for_defense.html.

¹⁷ See Tab 5 - Recommendation 2: “The Departments of State and Defense should negotiate formal AI cooperation agreements in the Indo-Pacific region with Australia, India, Japan, New Zealand, South Korea, and Vietnam” in *Interim Report and Third Quarter Recommendations*, NSCAI at 196 (October 2020), <https://www.nscai.gov/previous-reports/>.

¹⁸ *DoD Instruction 5530.03: International Agreements*, U.S. Department of Defense (Dec. 4, 2019), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/553003p.PDF>.

Chapter 5: AI and the Future of National Intelligence

Blueprint for Action

Intelligence will benefit from rapid adoption of artificial intelligence (AI)-enabled technologies more than any other national security mission. However, critical barriers keep the Intelligence Community (IC) from turning this potential into real capabilities that are scaled across agencies.

An Ambitious Agenda: AI-Ready by 2025.

To build on the progress that individual agencies have made, the IC should set the ambitious goal of adopting and integrating AI-enabled capabilities across every possible aspect of the intelligence enterprise as part of a larger vision for the future of intelligence.

Recommendation

Recommendation: Empower the IC's science and technology leadership.

Actions for Office of the Director of National Intelligence (ODNI):

- **The DNI should designate the Director of Science and Technology (S&T) as the IC Chief Technology Officer (CTO)¹ and direct the IC CTO to:**
 - o Develop and monitor IC-wide metrics for AI investments, AI implementation, AI outcomes, and AI readiness.
 - o Ensure maximum sharing and reuse of AI models, code, and tools across the IC to prevent unnecessary duplication where possible.
 - o Establish policies on, and supervise, IC research and engineering, technology development, technology transition, appropriate prototyping activities, experimentation, and developmental testing activities.
 - o After congressional approval and appropriation, manage a fund that would allow the ODNI to identify and invest in AI applications with outsized potential that may not have an identified source of agency or program funding as they near the end of their S&T life cycle.
- **The IC CTO, in coordination with the IC Chief Information Officer (CIO), Chief Data Officer (CDO), and Chief Information Security Officer, should oversee the establishment of common technical standards and policies for the IC. These standards and policies should be coordinated with the DoD to promote maximum interoperability, reciprocity, and data-sharing² in the following areas:**

- o An Application Programming Interface (API)–driven open architecture and associated policies that support the infrastructure to enable AI.³
- o Multi–level security standards for technical solutions allowing the movement of data across security clearance levels and the policies to enable it.
- o Data tagging and labeling.
- o Data sharing and access, including incentives for data stewards that reward their ability to share their data; shift the culture such that data stewards make it a default practice of externalizing their data via APIs, with appropriate levels of access restriction and control.
- o Common standards for machine readable processing, exploitation, and dissemination (PED) products.
- o Automated and reciprocal Authority to Operate (ATO) processes that include rapid code certification and accreditation processes.
- o Documentation strategies for data, models, and systems, and of the AI life cycle infrastructure to support traceability, training and testing procedures, and human–AI design guidelines.⁴
- o Technical standards for algorithms in support of interpretability and explanation, and policies to strengthen accountability.
- o Technologies and operational policies that align with privacy preservation, fairness, inclusion, human rights, and documentation of value considerations and trade–offs.⁵
- o Alternative hiring authorities for term–limited appointments appropriate for technical positions, such as Special Government Employees (SGE), highly qualified experts (HQE), and Intergovernmental Personnel Act (IPA) detailees.
- o Expanding the use of prize challenges as alternatives to traditional procurement.
- o Program and contracting guidance for well–documented and hardened APIs, data access and sharing across the IC, and provisions for the sharing and reuse of software products across the IC.
- **The IC CTO, in coordination with DoD, should develop a Technology Annex to the National Intelligence Strategy (NIS).⁶**
 - o The appendix should establish technology roadmaps to adopt AI–enabled applications to solve operational intelligence requirements. The appendix should address current issues within the IC, to include:
 - Aligning technical standards and policies with DoD to ensure seamless interoperability as well as make the Executive branch a better customer and more attractive market for industry.
 - Identify and promote acquisition reforms and methods that ensure the IC can rapidly procure and field systems to its intelligence professionals.
 - o The Technology Annex to the NIS should, at a minimum, include:
 - Intelligence support requirements, including how the IC analyzes the global environment and monitors technological advancements, adversarial capability development, scientific and technical cooperation among U.S. competitors, and emerging threats.

- Functional requirements and technical capabilities necessary to enable concepts that address each challenge.
 - A prioritized, time-phased plan for developing or acquiring such technical capabilities, that takes into account research and development timelines, a strategy for public private partnerships, and a strategy for connecting researchers to end users for early prototyping, experimentation, and iteration.
 - Additional or revised acquisition policies and workforce training requirements to enable IC personnel to identify, procure, integrate, and operate the technologies necessary to address the intelligence requirements.
 - Infrastructure requirements for developing and deploying technical capabilities, including data, compute, storage, and network needs; a resourced and prioritized plan for establishing such infrastructure; and an analysis of the testing, evaluation, verification, and validation (TEVV) requirements to support prototyping and experimentation and a resourced plan to implement them, including standards, test beds, and red-teams for testing AI systems against digital “denial & deception” attacks.
 - Consideration of human-factor elements associated with priority technical capabilities, including innovative human-centric approaches to user interface, human-machine teaming, and workflow integration.
 - Consideration of interoperability with allies and partners, including areas for sharing of data, tools, and intelligence products.
 - Flexibility to adapt and iterate appendix implementation at the speed of technological advancement.
- **ODNI should advance and continue to build out a purpose-built IC Information Technology Environment (ITE) that can fuse intelligence from different domains and sources.**
 - The IC ITE should be built in concert with the DoD digital ecosystem outlined in Chapter 2 of this report; they should focus on a federated system that is interoperable, integrated, and designed with building block services using the same interfaces as the DoD ecosystem.
 - The IC should accelerate ad hoc work and continuous experimentation to learn better how to integrate their systems.
 - Intelligence fusion promised by AI can only occur when all relevant data is made available across all systems. Building on the promise of IC ITE, the IC CIO and CTO should work with their counterparts across the IC and mission partners to ensure that IC integration and interoperability are given priority when evaluating technology investments.
 - The IC CTO should establish a multi-agency accredited learning environment and test bed where ad hoc work and continuous experimentation can occur using all relevant intelligence data.

Actions for Congress:

- **Designate the Director of S&T within ODNI as the IC CTO, and grant that position additional authorities for establishing policies on, and supervising, IC research**

and engineering, technology development, technology transition, appropriate prototyping activities, experimentation, and developmental testing activities.

- **Establish a fund that would allow the DNI to identify and invest in AI applications with outsized potential that may not have an identified source of agency or program funding as they near the end of their S&T life cycle.**
- **Grant the Director of National Intelligence sufficient budgetary authorities to enforce technical standards across the IC, including the ability to fence or otherwise withhold funding for programs that are not compliant with established common standards and policies.**
- **Establish a 10-year, \$1 billion, Program of Record to provide long-term, predictable funding for technologies identified in the Technology Annex to the National Intelligence Strategy.**
 - o This funding should target programs or departments with a proven track record of transitioning new or critical technologies to meet mission needs.

Recommendation: Change risk management practices to accelerate new technology adoption.

Recommendation

Actions for ODNI:

- **Establish an IT modernization Senior Risk Management Council (IT SRMC).**
 - o The IT SRMC should be tri-chaired by the IC CTO, CIO, and CDO to promote the effective governance of significant risk across the IC.
 - The IT SRMC should report to the Principal Deputy Director of National Security (PDDNI).
 - The IT SRMC should become a regular briefing entity in the Deputies Executive Committee (DEXCOM).
 - o The IT SRMC should include a senior member from the following IC entities:
 - ODNI Office of General Counsel
 - ODNI Office of Civil Liberties, Privacy, and Transparency
 - ODNI Mission Integration Directorate
 - Each intelligence agency and service branch
 - o The IT SRMC responsibilities should include:
 - Reviewing existing policies or creating new policies to ensure the IC uses informed risk acceptance and management practices when considering the adoption and use of AI technologies.
 - Advising the DNI on enterprise risk associated with not adopting AI technologies.
- **Address shortcomings in the current implementation of the National Institute of Standards & Technology (NIST) Federal Information Security Modernization Act (FISMA) Risk Management Framework (RMF).⁷**

- o Recommendations from the IT SMRC should inform the operational risk of not adopting a new technology as a balance to the technical risks considered in the RMF, allowing agencies to make better informed decisions on what systems they choose to bring on line or delay.
- o The IC should automate the implementation and simultaneous assessment of RMF considerations to the greatest extent possible.
- o Agencies within the IC often implement the RMF with different, but associated, policies that can prevent reciprocal accreditation and make it difficult to share tools among agencies.
- o The IC should make accreditation reciprocity within the RMF the standard and apply a high level of scrutiny to any agency that seeks to not recognize the accreditation of others.

Actions for Congress:

- **Assess the IC’s approach to risk and work with the IC to ensure the proper balance between risk acceptance, risk management, and risk avoidance.**

Recommendation

Recommendation: Improve coordination between the IC and DoD.

Actions for ODNI:

- **In coordination with the Secretary of Defense, the DNI should immediately issue a directive designating the PDDNI as a standing member and/or co-chair to the tri-chair Steering Committee on Emerging Technology.⁸**
 - o Absent of Congressional action, the Director of National Intelligence should work with the Secretary of Defense and members of the Steering Committee on Emerging Technology, including the Deputy Secretary of Defense and Under Secretary of Defense for Intelligence and Security, to identify the method and means to drive sustained coordination on emerging technology intelligence, policy, and resourcing.
- **Assist DoD, as requested, in developing the Technology Annex to the National Defense Strategy.⁹**
- **Work with DoD to establish an AI integration team focused on maximizing knowledge, data, and model sharing across and between the IC and DoD.**

Actions for Congress:

- **Revise the National Defense Authorization Act for Fiscal Year 2021 (FY2021 NDAA) provision authorizing a Steering Committee on Emerging Technology by designating it to be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.¹⁰**

Recommendation

Recommendation: Capitalize on AI-enabled analysis of open source and publicly available information.

Actions for ODNI:

- **Develop a coordinated and federated approach to integrate open source intelligence into all current intelligence processes and products. ODNI should promote coordination by taking the following actions:**
 - o Develop common standards and policies that enable the individual agencies to be more effective, such as contracting publicly available data sources for common use across the IC and clarifying or updating policy guidance on the appropriate use of publicly available and open source information, including with respect to privacy and civil liberties for U.S. persons or entities.
 - o Support the IC by identifying reliable industry partners across the spectrum of information sources and creating contract vehicles to rapidly integrate them into intelligence work across the IC. This should include establishing a pilot project to test “data-for-tools” exchanges in public-private partnerships.
 - o Aid the IC in communicating emerging risks and threats to industry and academia by coordinating the right expertise from across the IC; —for example, by connecting non-government entities to the Federal Bureau of Investigation for counterintelligence guidance, or to the U.S. Cyber Command/National Security Agency for cybersecurity.
 - o Develop a robust capability for bringing in individuals without security clearances or awaiting security clearance adjudication and allowing them to work on unclassified projects that directly support the IC.
- **Each individual agency should develop open source capabilities focused on the specialized applications of open source and publicly available information within their unique intelligence domains.**

Recommendation: Aggressively pursue security clearance reform for clearances at the Top Secret level and above, and enforce security clearance reciprocity among members of the IC.

Recommendation

Actions for ODNI:

- **Develop a Blueprint for Action for security clearance reform for clearances at the Top-Secret-and-above level including detailed timelines and metrics. The Blueprint for Action should include:**
 - o A collaborative effort with the private sector and academia to develop data-informed behavioral approaches to understanding risk factors and security clearance adjudication.¹¹
 - o Reforming identity management to ensure there is seamless security clearance reciprocity across the IC.
 - o A mechanism to enforce security clearance reciprocity among members of the IC and DoD.

Actions for Congress:

- **Congress should require the DNI to develop a Blueprint for Action for security clearance reform for clearances at the Top-Secret-and-above level including detailed timelines and metrics.**
- **Where necessary, Congress should reinforce the DNI's authority as head of the IC to enforce uniform security clearance policies and practices across the IC.**
- **Congress should require the DNI and the directors of the major intelligence services to regularly report on progress to the oversight committees.**

Blueprint for Action: Chapter 5 - Endnotes

¹ We envision the IC CTO as having roles, responsibilities, and authorities similar to the Under Secretary of Defense for Research and Engineering (USD (R&E)) within the DoD and to help the IC implement guidance and priorities established by the Steering Committee on Emerging Technology and the Technology Competitiveness Council.

² In Chapter 3 of this report, the Commission recommends the creation of a Steering Committee on Emerging Technology that is tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director for National Intelligence. This Committee should act as a forum through which to drive coordination between the IC and DoD, including the Chief Technology Officers.

³ Consistent with the DoD digital ecosystem described in the Chapter 2 Blueprint for Action, the API driven open architecture should: 1) Define a common set of well-documented common interfaces for the ecosystem's key components and building blocks; 2) Support and integrate the work of existing pathfinders up and down the ecosystem technology stack; and 3) Incorporate the process elements for data authorizations and continuous software ATO reciprocity.

⁴ Chapter 7 of this report provides more details on improving documentation practices for achieving baseline robust and reliable AI.

⁵ Chapter 8 of this report provides details on developing and testing systems per goals of privacy preservation and fairness.

⁶ A Technology Annex to the NIS should complement the Technology Annex to the National Defense Strategy (NDS) recommended in Chapter 2 of this report. The recommended Executive Agent for the Technology Annex to the NDS (see the Chapter 2 Blueprint for Action), the Under Secretary of Defense for Research and Engineering (USD (R&E)) should act as the primary interlocutor with the IC CTO for the creation of a Technology Annex to the NIS.

⁷ For more information, see *FISMA Implementation Project*, NIST (Dec. 3, 2020), <https://csrc.nist.gov/projects/risk-management/rmf-overview>.

⁸ The Chapter 3 Blueprint for Action calls for the Secretary of Defense, with support from the Director of National Intelligence, to issue a directive immediately establishing a tri-chair Steering Committee on Emerging Technology to oversee development of concepts and capabilities that include emerging and disruptive technologies to meet the current and future operational challenges facing the nation.

⁹ For a full discussion of the Technology Annex to the National Defense Strategy, see Chapter 2 of this report.

¹⁰ This action mirrors the Chapter 3 Blueprint for Action. While DoD and ODNI have the authority to establish such a forum without legislative action, codifying it into law will ensure that it is sustained through leadership transitions. If, at the drafting of the FY2022 NDAA, the DoD and ODNI have established the tri-chaired Steering Committee recommended herein, Congress should use the FY2022 NDAA to codify the body into law. If DoD and ODNI have not established the Committee as described in this report, Congress should include in the FY2022 NDAA a provision revising the FY2021 NDAA, section 236, which permits the creation of a Steering Committee on Emerging Technology, but is not structured effectively to improve coordination between the DoD and the IC. For a full discussion of section 236, see the Chapter 3 Blueprint for Action. The Commission also recommends that the legislative language be sufficiently broad so as to enable flexibility in the Steering Committee's roles and responsibilities should they need to adapt as emerging technologies and Department efforts evolve. See the Draft Legislative Language Appendix to this report.

¹¹ For more information on the need for an academic and scientific review of behavioral approaches to security clearance adjudication, see David Luckey, et al., *Assessing Continuous Evaluation Approaches for Insider Threats: How Can the Security Posture of the U.S. Departments and Agencies Be Improved?*, RAND Corporation at 28-34 (2019), https://www.rand.org/pubs/research_reports/RR2684.html.

Chapter 6: Technical Talent in Government

Blueprint for Action

The United States Government needs digital experts now or it will remain unprepared to buy, build, and use AI and its associated technologies. Expanding digital expertise is the most important step the government can take to modernize. While this challenge is recognized, few parts of government have adequately invested in building their digital workforce.

To expand its digital and AI digital workforce, the government needs to:

- *Organize* technologists within government through a talent management system designed to house highly skilled specialists.
- *Recruit* people that already have the skills the government needs, such as industry experts, academics, and recent college graduates.
- *Build* its own workforce by training and educating current and future government employees.
- *Employ* its digital workforce more effectively to ensure digital talent can perform meaningful work once they are in government.

Organize

Recommendation

Recommendation: Create Digital Corps for Cabinet-Level Departments and Select Agencies to Organize the Government's Technical Workforce

How a digital workforce is organized is as important as the workforce's level of expertise. We propose creating Digital Corps for Cabinet-level departments and select agencies that would recruit, train, and educate personnel; place personnel in and remove personnel from digital workforce billets; manage digital careers; and set standards for digital workforce qualifications. Agencies would create billets for members of the Digital Corps, and provide guidance to members of the Digital Corps about the work they perform.

Existing Models: The Army's Medical Corps. Full scaling of specialized talent will only happen if hired personnel have freedom to solve technical challenges. Many existing strategies for personnel management are inadequate due to a shortage of people in government agencies who can properly manage such specialized talent. A notable counterexample to this, which serves as an inspiration to our Digital Corps model, is the U.S. Army's Medical

Corps. The Medical Corps organizes experts with specialized health care skills that do not fit into the Army's traditional talent management framework.¹ Nurses and doctors receive education and training as civilians, but their skill sets are crucial to the Army's health care system. So, the Medical Corps talent management framework was created to house these medical professionals in a way that maximizes their ability to practice medicine within the Army. Like the Medical Corps, the Digital Corps should have specialized personnel policies, guidelines for promotion, training resources, and certifications for personnel to demonstrate proficiency in new digital areas.

Notably, a Digital Corps would not be comparable to either the Marine Corps or a Space Service, as it would not have a service secretary or a distinct theater or domain, and its members would work for existing services or agencies.

Roles Within the Digital Corps. Career fields are distinct from core competencies—skills that every Digital Corps member should possess prior to hiring—such as modern stack software development, deployment, and data-informed decision-making. Training resources for each career field should be made available to Digital Corps members across every agency. Departments and agencies must also be cognizant that digital talent is rarely interchangeable across different skill sets; for example, database architecture, machine learning, and user experience design all fall into different career fields with near-zero overlap. Digital Corps members should be allowed to focus on any one of the following additional career fields:

- Software development
- Data science
- Artificial intelligence
- DevOps and site reliability engineering
- Human-centered product design
- Product management
- Security
- Data governance and use
- Emerging technologies²

Digital Corps technologists should be able to continue to promote without leaving their focus area and move upward into management. Many private tech companies distinguish between their engineering and engineering management tracks, so that skilled engineers are not incentivized to become managers solely for the sake of career advancement. The Army's Medical Corps follows a similar model. Once promoted, officers highly competent in their medical specialty can either continue as clinicians or become administrators and managers within the Medical Service Corps.

Staffing and Digital Corps Billets. Cabinet-level departments and select agencies should develop their own Digital Corps rather than relying on a single, government-wide Digital Corps. For Corps members, this approach creates well-defined tracks for career progression and stronger incentives to stay. This approach also makes it easier for departments and agencies to identify and invest in in-house talent for future technology projects.

Each Cabinet-level department and select agency should create designated billets to be filled by qualified members of its Digital Corps based on skills and experience. In addition, each should maintain a central talent repository with Corps members' portfolios of prior digital projects completed with the agency. Departments and agencies can then search this repository to find the most suitable Corps member to fill each billet. Taking inspiration from software development companies, one method of reliably measuring skill proficiency is to conduct digital interviews consisting of case questions and whiteboarding exercises. We recommend that billets be filled based on candidates' performance in these interviews, chosen career field, and prior project experience (possibly while filling other billets within the same agency at an earlier date).

Actions for Departments and Select Agencies:

- **Allocate resources toward the creation of Digital Corps modeled after the Army's Medical Corps.**
- **Develop Digital Corps training resources in the forms of licensed instructional videos, tutorials, and coursework for each of the nine career fields listed.**
- **Create agency-specific talent repositories where Corps members can list project portfolios, source code (where permitted), and career field training badges.**
- **Create billets and fill them through interviews, evaluation of Corps members' career field training, and other relevant experiences.**
- **Develop parallel management-oriented and technical-oriented tracks for each Corps member's career progression, with set standards for promotion per agency.**

Recruit

The government needs to improve its ability to attract scarce AI talent from the private sector, academia, and recent college graduates. Doing so requires making paths to service as easy as possible for as many technologists as possible.

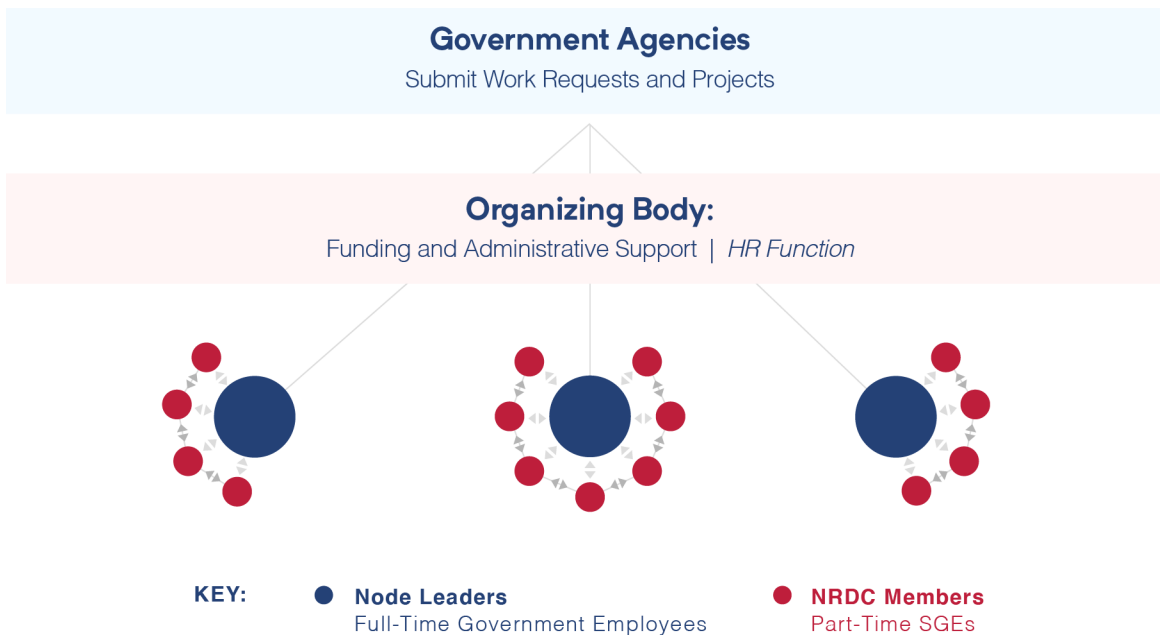
Many AI and other digital practitioners are interested in working with the government and can and would do so as either full-time employees or part-time employees. Of those desiring full-time employment, some seek an entire career as a government civilian or in the military. Others, while willing to work with the U.S. Government full-time, are less willing to make long-term commitments or to dedicate as much of their time, and instead desire to become short-term employees, fellows, talent exchange participants, or military reservists. A third group is willing to work with or for the government part-time, but are unwilling to become full-time civilian employees and have no desire to serve as part of the military. To improve recruiting, the government needs to improve the hiring process and build mechanisms for part-time civilian service.

Recommendation: Create a National Reserve Digital Corps

Recommendation

The government would benefit from access to a larger portion of the country’s total digital workforce. Many government digital projects suffer from lack of access to digital expertise. The U.S. Government should establish a civilian National Reserve Digital Corps (NRDC) modeled after the military reserves’ service commitments and incentive structure. Members of the NRDC would become civilian special government employees (SGEs),³ and work at least 38 days each year as short-term advisors, instructors, or developers across the government.⁴ Longer-term positions would be established on an individual basis. While short-term volunteers are not a substitute for full-time employees, they can help improve AI education for both technologists and non-technical leaders, perform data triage and acquisition, help guide projects and frame technical solutions, build bridges between the public and private sector, and other important tasks.⁵ Several AI practitioners within the United States Government have said during interviews with the NSCAI that their projects would benefit from the kind of reserve corps we propose here.

National Reserve Digital Corps.



General Structure. We recommend establishing and managing the NRDC as a set of nodes that fall under the supervision of the Office of Management and Budget (OMB). Each node would be aligned with a full-time government employee leader selected by OMB rather than geography, digital applications, or government agency. In effect, OMB would select node leaders, who would then be responsible for recruiting and organizing their team. In addition to selecting node leaders, OMB would establish standards, ensure nodes meet government client requirements, provide funding and administrative support, maintain security clearances, establish access to an agile development environment and tools, and facilitate technical exchange meetings, when appropriate, to ensure stovepipes are not created.

Recruitment. Each node would be responsible for recruiting and screening its digital experts. Notably, OMB would not be responsible for establishing qualification standards for members of the NRDC. While volunteers would need to be able to pass a background check and would not be employees of a foreign government (though they might be foreign nationals), node leaders would be empowered to screen and select volunteers, and to recruit experts from within NRDC for specific tasks. OMB would provide administrative support, much like a human resources team in a private sector company.⁶

Project Selection. Projects would be selected in three ways:

- **Selection by a node after contact with a government client**
- **OMB would direct a node to take on a project**
- **Node leadership would approve individual projects driven by a perceived need that is not tied to a request from a government client**

Government clients would directly contact node leaders or OMB. Nodes would be responsible for establishing relationships with government agencies and selecting projects, but OMB would be responsible for ensuring that agencies' requests are received and that nodes contribute to NRDC's mission and vision. Individual projects that are not driven by a government client's request would be pursued at the node leadership's discretion.

Relationship with Government Agencies. Members of the NRDC would work with agencies on a project-to-project basis, such as consulting for a specific project or teaching a specific course. They would not have a commitment to work with the same agency consistently. Government agencies would be responsible for paying for their projects, including the cost for reservist time.

Relationship with Civilian Employers. Members of the NRDC and their civilian employers would be bound by the same rules as the military reserve under the Uniformed Services Employment and Reemployment Rights Act (USERRA).⁷ Members would be responsible for identifying conflicts of interest and removing themselves as appropriate. Employers would not be able to discriminate against members of NRDC, fire them, or delay promotions as a consequence of spending time serving in NRDC.⁸ Implementation could take the form of a legislative recommendation to modify USERRA or a proposal modeled after USERRA.

Incentivizing Reservist Participation. Civilian reservists in this program would benefit in several ways. They would gain an opportunity to contribute to their country, do exciting, meaningful work, and attain awareness of work and advances in a community that differs from their own. They may also benefit from the following incentives:

- **The government should create an NRDC scholarship program modeled after the Reserve Officer Training Corps. Universities would select students through a competitive process to receive full tuition and study specific disciplines related to digital technology. In return for accepting the scholarship, graduates would spend part of their summers during school in government internships. Between their freshman**

and sophomore years, students would spend six weeks becoming familiar with a range of U.S. Government departments and agencies. Between their sophomore and junior years, students would spend six weeks as an intern at a specific government agency or office. Between their junior and senior years, students would spend another six weeks interning at a specific agency or office. Upon graduation, scholarship recipients would spend five years serving in the NRDC, beginning as a GS-7 and advancing to a GS-11 over the course of five years. Students would also begin the security clearance process at least two years before graduating.⁹

- The NRDC should include a training and continuing education fund for all members. The NRDC would pay up to \$50,000 to each reservist to attend training and educational opportunities related to AI or to pay for student loans. Educational opportunities would include conferences, seminars, degree and certificate granting programs, and other opportunities. An incentive explicitly tied to continuing education would increase the perceived and actual competency of AI reservists. It would also attract those with an active interest in continuing education, especially new practitioners seeking to establish themselves.

How NRDC Would Work: An Example. The following is a hypothetical example of how the NRDC would function. In this example, OMB would begin creating a node by selecting a leader that would be trusted to establish and manage a team of reservists. OMB selects “Jennifer,” a full-time government employee working within the NRDC division of OMB, to lead a new NRDC node. Jennifer decides to organize her node functionally rather than regionally. Using existing government tools and her professional contacts, she recruits people from across the country, most of whom have backgrounds in health care data management or recent graduates with degrees related to the field. She also recruits from within the NRDC by posting open positions on online job boards. During the recruitment process, OMB provides financial support for recruitment efforts, travel money, and processes new reservist administrative paperwork, including security clearance applications.

After the node is established and the team is in place, a government agency—in this example, the Centers for Disease Control and Prevention (CDC)—realizes it has two digital needs it cannot meet internally: improving a database and training their workforce in new data management practices at the National Center for Chronic Disease Prevention and Health Promotion. After reaching out to OMB, they determine that Jennifer’s node is the best fit, and request assistance. After examining the request and her team’s workload, Jennifer determines that she would support the CDC’s database improvement request with a four-person team and support workforce training with a two-person team. The four-person team spends 14 days examining the existing database and making updates to the database. The two-person team spends 10 days on site at the National Center for Chronic Disease Prevention and Health Promotion speaking with leaders and employees about their data management needs and the current state of the workforce’s skill level, developing curriculum, and teaching data management best practices.

The teams Jennifer selects to support the CDC include Michael. Michael received a four-year scholarship from NRDC to study computer science as an undergraduate. After graduating three years ago, he began working full-time as a data analyst at a health care company and working part-time on NRDC projects he coordinates with his node leader. He also used his education stipend to pay for an online course from MIT last year. This hypothetical shows that an NRDC can effectively increase the U.S. digital talent, connect private-sector workers with a government agency, and create a pathway for that connection to solve an actual problem.

Actions for Congress:

- **Pass legislation establishing the NRDC within OMB**

- o Grant OMB direct-hire authorities to hire node leaders and reservists.
- o The NRDC should offer full tuition scholarships to students studying specific disciplines related to national security digital technology for up to four years in exchange for five years of service as a member of the NRDC. This could be done by including service in the NRDC as an option for people with degrees in digital fields to pay off service obligations incurred as a result of education received in the Defense Civilian Training Corps.¹⁰
- o Legislation should authorize up to \$50,000 in educational benefits for courses, seminars, conferences, and other educational opportunities that are approved by OMB. It should also ensure that members of the NRDC receive the same employment protections as military reservists under USERRA. This can be done by amending USERRA to cover “service in the uniformed services or the National Reserve Digital Corps.”
- o Congress should make a two-year appropriation of \$16 million to pay for initial administrative, scholarship, and education benefits.

- **Evaluate NRDC Success**

- o Use three metrics to evaluate NRDC’s success: 1) The number of technologists who participate annually; 2) Evaluations of results from government clients; and 3) Evaluations of results from reservists. OMB should establish the central, organizing function for the NRDC within six months of the passage of legislation, and establish five nodes and a mechanism for distributing educational benefits within nine months of the passage of legislation.

Actions for OMB:

- **Immediately upon receiving authority from Congress, establish a National Reserve Digital Corps with systems and processes designed to:**

- o Select and hire node leaders
- o Encourage potential government clients to contact NRDC nodes, or OMB, with potential problems to resolve
- o Ensure government client needs are met by NRDC nodes
- o Provide funding for education supplements and scholarship programs

- o Provide administrative support (including for security clearances)
- o Establish node access to development environments and tools
- o Facilitate technical exchange meetings
- o Match recipients of NRDC scholarships with node leaders
- **At the outset, establish five NRDC nodes. Each node leader should be responsible for:**
 - o Recruiting and hiring reservists
 - o Ensuring the quality of their work
 - o Partnering with government agencies

Recommendation: Create Digital Talent Recruiting Offices Aligned with Digital Corps

Recommendation

Executive branch agencies should create agency-level digital talent offices of up to 20 personnel responsible for recruiting both early career and experienced professionals. Recruiting offices would monitor their agencies' need for specific types of digital talent. The offices would be empowered to recruit technologists virtually, by attending conferences, career fairs, recruiting on college campuses, and offering scholarships, recruiting bonuses, referral bonuses, non-traditional recruiting techniques such as prize competitions, and other recruiting mechanisms. A recruiting office would assume responsibility for their agency's digital talent recruitment efforts, e.g., Science, Mathematics and Research for Transformation (SMART) Scholarship-for-Service, and partner with agency human resources offices to use direct-hire authorities and the Intergovernmental Personnel Act (IPA) to accelerate hiring. This would help scale digital talent recruitment by creating a central, empowered organization that focuses on a specific mission; concentrates expertise and funds; would help experts move in and out of government positions throughout their career; and can develop relationships with universities and private-sector companies.

Actions for Congress:

- **Amend Section 230 of the FY2020 NDAA. (Armed Services Committees)**
 - o The DoD should be required to appoint a civilian official responsible for digital engineering talent recruitment policies and their implementation.
 - o The civilian official should be supported by a digital talent recruiting office with the Office of the Under Secretary for Personnel and Readiness, as described above.
- **Require the Office of the Director of National Intelligence (ODNI) to create a digital talent recruiting office. (Intelligence Committees)**
 - o The office should work with the IC to identify their agencies' needs for specific types of digital talent.
 - o Recruit technologists by attending conferences, career fairs, and actively recruiting on college campuses.

- o Integrate federal scholarship for service programs into agency recruiting; offer recruitment and referral bonuses.
- o Partner with their agencies' human resource teams to use direct-hire authorities to accelerate hiring.
- **Require the Department of Homeland Security (DHS) to create a digital talent recruiting office. (Senate Homeland Security and Governmental Affairs Committee and the House Committee on Homeland Security)**
 - o The office should work with DHS to identify their agencies' needs for specific types of digital talent.
 - o Recruit technologists by attending conferences, career fairs, and actively recruiting on college campuses.
 - o Integrate federal scholarship for service programs into agency recruiting; offer recruitment and referral bonuses.
 - o Partner with their agencies' human resource teams to use direct-hire authorities to accelerate hiring.
- **Require the Department of Energy (DoE) to create a digital talent recruiting office. (Senate Committee on Energy and Natural Resources and the House Committee on Energy and Commerce)**
 - o The office should work with DoE to identify their agencies' needs for specific types of digital talent.
 - o Recruit technologists by attending conferences, career fairs, and actively recruiting on college campuses.
 - o Integrate federal scholarship for service programs into agency recruiting; offer recruitment and referral bonuses.
 - o Partner with their agencies' human resource teams to use direct-hire authorities to accelerate hiring.

Actions for DoD, including U.S. military services, DOE, DHS, and the ODNI:

- **Create digital talent recruiting offices.**
 - o Offices should work with their agencies to identify their need for specific types of digital talent.
 - o Recruit technologists by attending conferences, career fairs, and actively recruiting on college campuses.
 - o Integrate federal scholarship for service programs into agency recruiting; offer recruitment and referral bonuses.
 - o Partner with their agencies' human resource teams to use direct-hire authorities to accelerate hiring.

Recommendation: Grant exemption from OPM General Schedule Qualification Policies for Specific Billets and Position Descriptions

Recommendation

AI practitioners applying for positions within the federal government and their hiring agencies are constrained by OPM minimum qualification standards. While these standards are important, and have increased fairness in hiring, they also prevent expert technologists that do not have master's degrees—and in some cases, bachelor's degrees or comparable work experience—from joining the government at a reasonable level of compensation. For example, a 19-year-old software developer or AI practitioner might have a proven track record on cybersecurity or in AI competitions, but can only enter the government as a GS-7. To reduce this hiring challenge, the government should allow agencies to exempt certain billets from OPM general schedule qualification policies, and instead allow local hiring managers to make an independent decision about both hiring and pay grade based on evaluations, prior work, alternative certification programs, or practical experience.

Actions for Congress:

- **Direct the Office of Personnel Management to amend 5 C.F.R. § 338.301, on service appointments.**
 - Allow service secretaries and cabinet officials to create exceptions from the Qualification Standards for General Schedule Positions by individual billet or position description.

Actions for OPM and Military Services:

- **OPM should create and execute a process by which federal departments and agencies can apply for billets or position descriptions to be exempt from general schedule qualification policies.**
- **Two-star-and-above commands and their civilian equivalents should declare individual billets and position descriptions exempt from OPM qualification standards without approval from OPM or any more senior authority.**

Recommendation: Expand the CyberCorps: Scholarship for Service

Recommendation

The CyberCorps: Scholarship for Service (SFS) is a recruiting program designed to attract students studying IT, cybersecurity, and related fields into the USG. Expanding it could bring in more people with AI-related skills. It is managed by the National Science Foundation in partnership with the Office of Personnel Management and the Department of Homeland Security. Students enrolled in the program receive a scholarship in exchange for an obligation to work in an approved government agency for a period of time equal to the time covered by the scholarship. Seventy undergraduate and graduate institutions participate in SFS by selecting students for the program, and since 2001, 3,600 students have received scholarships, 94% of whom went on to serve in government.¹¹ Hiring typically takes place during annual online and in-person career fairs.¹²

It should be noted that cyber and AI are different fields. Expanding CyberCorps: SFS to CyberCorps and AI: SFS would avoid increasing administrative burdens. This should not be taken as an indication that AI and cyber are synonymous, as the education and skills for each field differ.

Actions for Congress:

- **Amend the CyberCorps: SFS, as defined by Section 230 of the National Defense Authorization Act for Fiscal Year 2020.**
 - Include digital engineers.
 - Pay for up to four years of scholarships.
 - Include the opportunity to begin the security clearance process.
- **Amend 15 U.S.C. § 7442 subsection (a).**
 - ... recruit and train the next generation of information technology professionals, digital engineers, artificial intelligence practitioners, data engineers, data analysts, data scientists, industrial control system security professionals, security managers, and cybersecurity course instructors to meet the needs of the cybersecurity mission for Federal, State, local, tribal, and territorial governments.
- **Amend 15 U.S.C. § 7442 subsection (b).**
 - Provide an opportunity for scholarship recipients to initiate their security clearance process at least one year before their planned graduation date.
- **Amend 15 U.S.C. § 7442 subsection (c).**
 - Allow the scholarship to last for 4 years.

Actions for the National Science Foundation and Office of Personnel Management:

- **Broaden the CyberCorps: SFS.**
 - Pay for up to four years.
 - Include fields falling under digital engineering, as those fields are defined by the National Defense Authorization Act for Fiscal Year 2020 (Pub. L. 116–92, section 230): the discipline and set of skills involved in the creation, processing, transmission, integration, and storage of digital data, including data science, machine learning, software engineering, software product management, and artificial intelligence product management.

Recommendation

Recommendation: Establish a STEM Corps

A bipartisan group of members of the House Armed Services Committee have proposed H.R. 6526, STEM Corps Act of 2020. The proposal would authorize the appropriation of \$5 million per fiscal year, with \$500,000 for administrative costs and an advisory board. The program provides a maximum scholarship of \$40,000 per student per year. Scholarship recipients would serve in different capacities within the DoD for a minimum of three years,

with an option to either remain in the DoD or transfer to a private-sector company that has contributed to STEM Corps funding. The proposal requires participants to be paid at a rate not less than GS-6 for the first three years of their obligation and at not less than as a GS-10 during their fourth year. This proposal has the potential to significantly increase the number of personnel with STEM backgrounds in the DoD civilian workforce for a relatively low cost if a sufficient number of private-sector companies contribute. The potential for recipients to transfer to the private sector after three years of government service may create retention issues, but it may also serve as a mechanism to create bridges between the DoD and private sector companies.

Actions for Congress:

- **Establish a STEM Corps in the FY2022 NDAA.**
- **Set aside \$5 million for a STEM Corps for FY2022 and each fiscal year thereafter.**

Actions for the DoD:

- **With congressional authorization and appropriation, establish an office to manage and establish a STEM Corps as described above.**
- **Include a scholarship program, advisory board, private-sector partnership program, and STEM Corps member management program.**

Build

The government will not be able to come out of its workforce deficit through recruiting alone. AI and digital talent is simply too scarce in the United States. In 2020, there were more than 430,000 open computer science jobs in the United States, while only 71,000 new computer scientists graduate from American universities each year.¹³ To overcome the challenges presented by AI and digital talent scarcity, the government should deliberately focus on building its AI and digital workforce.

Recommendation: Create a United States Digital Service Academy

Recommendation

The United States needs a new academy to train future public servants in digital skills. Civil servants play a critical and often underappreciated role in government. They hold much of the government's niche, long-term expertise. This is especially true for the digital expertise that is badly needed for the government to modernize. Methods like the competitive service and scholarship for service programs have helped recruit talent, but as the government's needs changed, those approaches will no longer address the full scope of the government's needs. Bolder measures are necessary to produce the broad, diverse, and technically educated workforce the government needs.

Our proposed United States Digital Service Academy (USDSA) would be an accredited, degree-granting university that receives government funding,¹⁴ be an independent entity

within the federal government, and have the mission to help meet the government's needs for digital expertise. It would be advised by an interagency board that would be assisted by a federal advisory committee composed of commercial and academic leaders in emerging technology.

Existing Models: The Military Service Academies. The USDSA should be modeled off of the five U.S. military service academies but should produce trained government civilians not only to the military departments, but also to civilian departments and agencies beyond DoD.¹⁵

The five military service academies each produce commissioned officers for the armed forces.¹⁶ The academies select cadets and midshipmen through a congressional and presidential nomination process, followed by a competitive admissions process. The cadets and midshipmen, who are government employees, exchange a commitment to serve after graduation for a tuition-free education. Many choose this path for the opportunity to serve; the free tuition and education often are considered a bonus. Those who depart prior to meeting the minimum requirements for graduation still incur either a service commitment or financial requirement to pay back education received upon their departure from the schools.

The academies contribute between 15% and 20% of the new junior officers to their respective services each year—the largest single commissioning source. Academy graduates also play an outsized role in the military services' senior leadership.¹⁷ As a result, the academies help shape the identity and culture of their services, including their standards and ethical norms. USDSA would be comparable to the other service academies in many ways. It would be a degree-granting institution focused on producing leaders for the United States Government. USDSA students, like military service academy students, would not pay for tuition, or room and board, and would have a post-graduation service obligation. Americans should expect USDSA graduates to seek to serve, to lead the nation's digital workforce, and to ensure the United States sets an example of intelligent, responsible, and ethical high-tech leadership.

Key Differences Between USDSA and the Military Service Academies. The USDSA would differ in significant ways. First and foremost, USDSA students would enter the institution to become civil servants. They would know that their education would be repaid in the form of a five-year obligation to serve in government, which would begin upon graduation when they become a civil servant at a GS-7 level. Exclusively producing civil servants would eliminate the need for students to complete commissioning requirements, simplifying the school's curriculum and administrative burdens, and reduce the need for expansive campus real estate for training and parade grounds. It would also make USDSA less redundant, as the military service academies already produce hundreds of computer scientists, electrical engineers, and mathematicians every year.

USDSA students would also have a more STEM-focused education. While the core curriculum would ensure broad exposure to different fields, students would have a highly technical education. A wide variety of technical majors could include AI, software engineering, electrical science and engineering, computer science, molecular biology, computational biology, biological engineering, cybersecurity, data science, mathematics, physics, human-computer interaction, robotics, and design. Students could also blend those majors with humanities and social science disciplines such as political science, economics, ethics and philosophy, or history.

A third difference would be that USDSA graduates would serve across the Federal government. To avoid both perceived and real parochial bias from the organizations that administer service academies, USDSA would be administered as an independent Federal entity. The minimum and maximum number of graduates who would serve in each department or agency would be determined annually by an interagency board.¹⁸

Mission Statement of the USDSA. We propose the following: “The United States Digital Service Academy’s mission is to develop, educate, train, and inspire digital technology leaders and innovators and imbue them with the highest ideals of duty, honor, and service to the United States of America in order to prepare them to lead in service to our nation.”

The Student Experience. During their first year, students would begin the Academy’s core curriculum, explore some electives to help determine their major, and take a summer internship or fellowship. The core curriculum is envisioned to include, among other things, American history, government, and law, as well as composition, mathematics, computer science, and the physical and biological sciences. Once summer arrives, students would participate in summer internships with private sector companies.

Students would select their major early in their second year, begin concentrating on their technical field, and continue their core curriculum. They would also initiate their security clearance application process. The goal would be for all students to graduate with at least a secret clearance. After completing the classroom portion of their second year, students would complete internships in two government agencies, which would help them focus their goals for government service.

During their third year, USDSA students would increase the focus on their major, complete the majority of their core curriculum, and begin committing to a government agency. Similar to the military service academies, attendance of the first day of class at the start of their third year serves as a commitment to five years of government service upon graduation. After completing the classroom portion of the third year, students would participate in another private sector internship.

Students would commit to a particular government agency and career field during the first weeks of their fourth year and begin the job placement process. To select a

department and career field, students would create a rank ordered list of career fields within departments, agencies, and services. The USDSA would then match student preferences to the government's needs as identified by an annual interagency process. After successfully completing all academic requirements, students would graduate as GS-7s, with the potential to progress rapidly to GS-11. After completing their service obligation, USDSA graduates would have the opportunity to transition to the NRDC.

Accreditation. In order to receive federal funding, the USDSA would take the required steps to complete the accreditation process through a regional accreditation organization. The accreditation organization would be determined based on the physical location of the institution and recognized by the Department of Education and Council for Higher Education Accreditation.¹⁹ Membership in such an organization ensures academic quality throughout the institution's life span, as accreditation requires ongoing assessment for improvement. Future employers are able to affirm the credentials of USDSA graduates, the academy is able to accept charitable donations, and post-graduate programs recognize the validity of undergraduate degrees through accreditation. Based on the location of USDSA, the institution would also work with the hosting state to determine compliance with all core standards and processes.²⁰

Proposed Blueprint for Action for the USDSA:

Phase One (Years 1-2)

- Identify and secure an appropriate site for initial USDSA buildout with room for future expansion.
- Identify gaps in the government's current and envisioned digital workforce by an interagency task force under Office of Personnel Management leadership.
- Establish the USDSA administration as a new Executive branch agency with an individual appropriation that will be responsible for the phased Blueprint for Action plan and the management of the institution.
- Recruit tenure-track faculty.
- Recruit adjunct faculty, primarily from private-sector technology companies.²¹
- Grant the USDSA the authority to accept outside funds and gifts from individuals and corporations for startup, maintenance, and infrastructure costs.
- Appropriate \$40 million to pay for administrative costs.
- Satisfy the necessary requirements set by the Department of Education as well as the state USDSA is in for degree-granting approval.
- Apply for degree-program-specific accreditation through Computing Accreditation Commission on Colleges of Accreditation Board for Engineering and Technology.²²
- Apply for accreditation with a Regional Accrediting Organization approved by the Department of Education and Council for Higher Education Accreditation in order to be granted "Candidate" status.

- Construct initial physical infrastructure.
- Appropriate additional costs for the selection and purchase of the physical location and construction of infrastructure.

Phase Two (Years 3-5)

- Begin classes with an initial class of 500 students at the beginning of year three.²³
- Demonstrate compliance with all requirements and standards of the regional accrediting organization in order to be granted Membership status.

Phase Three (Years 6-7)

- Graduate the first class.
- Ongoing improvement through accreditation assessments.
- Assess, and as appropriate, expand class sizes.

Actions for Congress:

- **Authorize the establishment of the USDSA.**
 - An independent entity with a mandate to establish the institution described above.
 - Appropriate \$40 million over two years to pay for the USDSA's initial administrative costs.

Actions for the Office of Personnel Management:

- **Begin an interagency process to identify skill and personnel gaps in the federal government's digital workforce.**

Employ

Digitally talented people should be able to reasonably expect to spend a career performing meaningful work focused on their field of expertise in government. Without such an expectation, they are unlikely to join the government workforce, and without their experience matching expectations, they are unlikely to stay for long.

Recommendation: Establish Career Fields for Government Civilians in Software Development, Software Engineering, Data Science, Knowledge Management, and Artificial Intelligence

Recommendation

Government civilians play a critical role in the national security enterprise. A significant portion of the government's AI talent is likely to exist in the civilian workforce. Government civilians currently do not have career paths outside of research and development that

allow them to focus on software development, data science, or AI for the majority of their career. This results in a highly limited ability to recruit talent from outside of government, an inability for an individual to focus on a skill set for an extended time, a lack of continuing education opportunities for these government civilians, and retention issues. It also causes the government to struggle to identify and manage the software development, data science, and AI talent within its workforce.²⁴ Digitally focused occupational series will better allow the government to track and manage its digital workforce, to attract new talent that wants to focus on a technical skill set, and to create new positions.

The government should create software development, software engineering, data science, knowledge management, and AI occupational series. This combination of occupational series would significantly improve the government's ability to recruit and manage experts that will supervise the collection and curation of data, build human-machine interfaces, and help end users generate and act on data-informed insights. Many successful private-sector organizations use a version of this combination of skills.²⁵ The government should follow their example.

Actions for Congress:

- **Require OPM to draft software development, software engineering, data science, knowledge management, and artificial intelligence occupational series classification policies no later than 270 days after the passage of the legislation.**

Actions for OPM:

- **Create software development, software engineering, data science, knowledge management, and artificial intelligence occupational series.**
- **Accelerate the creation of new digital occupational series.**
 - o Rather than waiting for agencies to provide a formal request for a new occupational series, ask agencies to provide supporting documents and subject matter experts to study and draft a classification policy for each occupational series.

Recommendation

Recommendation: Establish Digital Career Fields for Military Personnel

Digital subject matter experts' inability to spend a career working on digital topics while serving in the military is arguably the single most important issue impeding military modernization.²⁶ Much like their civilian counterparts, U.S. military personnel do not have career paths that allow them to focus on software development, data science, or AI for the majority of their career.²⁷ The military has established career fields for doctors and lawyers that allow them to focus on a technical field, develop their skill over time, and advance within their service. The military is choosing not to do the same for many types of digital talent. While some of the services train some operational research and systems analysis (ORSA) personnel to perform machine learning and AI tasks, these personnel may be shifted to work on other ORSA tasks rather than AI. Phrased differently, AI practitioners have some background in ORSA, but not all ORSA personnel are trained to work in machine learning or AI.²⁸

This results in a reduced ability to recruit talent outside of the government, an inability to focus on a skill set for an extended time, a lack of continuing education opportunities, and retention issues. It also causes the government to struggle to identify and manage the software development, data science, and AI talent within its workforce.²⁹ These problems are particularly acute for military personnel, who are required to regularly change positions and move into manager roles or face eventual discharge from the military. The lack of digital career fields also causes the military services to struggle to identify and manage the software development, data science, and AI talent within their workforces.³⁰ As long as this state continues, the military should not expect to achieve better results for its digital modernization than its legal and medical fields would have without career fields for lawyers and doctors.

The military services should have primary career fields that allow military personnel to focus on software development, data science, or artificial intelligence for their entire career, either as managers or technical specialists.

Actions for Congress:

- **Require the military service chiefs to create career fields focused on software development, data science, and artificial intelligence.**
 - o Congress should amend section 230 of the FY2020 NDAA to require the military service chiefs to create career fields focused on software development, career fields focused on data science, and career fields focused on artificial intelligence for both commissioned officers and enlisted personnel, and, as appropriate, warrant officers.
 - o Military personnel should be able to join these career fields either upon entry into the military, or by transferring into the field after serving a period in another career field. These career fields should have options that allow personnel to either follow a path to senior leadership positions, or specialize and focus on technical skill sets. Those that specialize and focus on technical skill sets should not have to leave their focus area and move into management positions to continue to promote. Legislation should not restrict the military services to only two career fields, but rather require each service to create at least two career fields, and more at their discretion. The military services should be required to create the career fields within one year of passage of legislation.

Actions for the Military Services:

- **Create career fields that allow military personnel to focus on software development, career fields that allow military personnel to focus on data science, and career fields that allow military personnel to focus on artificial intelligence.**
 - o While remaining consistent with service personnel policies and procedures, these career fields should be open to both enlisted personnel and commissioned officers, and, as appropriate, warrant officers.
 - o Military personnel should be able to join these career fields either upon entry into the military, or by transferring into the field after serving a period in another career field.

- o These career fields should have options that allow personnel to either follow a path to senior leadership positions, or specialize and focus on technical skill sets. Those that specialize and focus on technical skill sets should not have to leave their focus area and move into management positions to continue to promote.

Recommendation

Recommendation: Provide Government Technologists with World-Class Tools, Data Sets, and Infrastructure.

Highly skilled technologists working in government are regularly denied access to software engineering tools. They have to jump bureaucratic hurdles to accomplish basic job functions such as sharing source code or downloading data sets, leading to frustration and periods of idling. To perform meaningful work in government, employees within the digital workforce need access to enterprise-level software capabilities at par with those found in the private sector. Capabilities include software engineering tools, access to software libraries, open-source support, and infrastructure for large-scale collaboration. Employees within the AI career field in particular will need access to further specialized resources such as curated data sets and compute power.

In order to be effective, developers need to be able to find and view source code written by other developers before them. Being unaware of existing code repositories often leads to writing redundant software that meets a different set of quality standards and robustness than existing software. To prevent this, each member of the AI career field needs access to a shared, enterprise-level repository of AI software and tools, similar to that recommended in Chapter 2 of this report for the Department of Defense. This repository should house source code available to all AI developers within a government agency.

Each government agency should create enterprise-scale solutions for source code management across multiple software projects. This does not mean that every developer in an agency will be able to modify every single project in a repository—with protocols for delegated access, a system administrator can set project-specific read and write permissions for each AI developer. New software projects should be set up to allow ubiquitous unit testing as code is written, and automatic integration into a code review process to ensure robust and bug-free output. Following these guidelines will promote a culture of software engineering excellence, emphasizing to technologists that it is possible to work in government while remaining at the forefront of a digital field.

For new developers who join an agency, onboarding procedures must include separate instructions for pushing their new code to this repository as well as instructions on how to navigate the software catalog and search for existing source code.

All career fields also need unobstructed access to the latest open-source libraries and tools. Over time, technologists develop individual preferences for their software development environment, opting for custom software development kits (SDKs), debugging tools, cloud tools, version control software, and data visualization platforms on local machines. To

ensure productivity and developer satisfaction, agencies must give each developer the authority to install vetted, authorized tools on their local machines.

AI developers use open-source software libraries for training machine learning models and making them production-ready for real-world use. To harness the full power of these essential libraries, AI developers should have access to vetted libraries, but also to compute power while training their machine learning models. Models train very slowly on a local machine because of the complexity of underlying mathematical calculations in the training process. As a result, AI developers prefer to train them rapidly through automatic deployment pipelines on commercially available platforms, or another external service. Smoothing the transition from local software development to cloud services is critical for any organization using AI and ML.³¹

Actions for Departments and Agencies (including, but not limited to, the Department of Energy, Department of Homeland Security, Department of State, Department of Commerce, and Department of Justice):³²

- **Ensure software developers and engineers, data scientists, and AI practitioners:**

- o Have access to systems with capabilities comparable to Repo One and Platform One.
- o Are authorized to install custom software licenses, debugging tools, cloud deployment tools, version control software, and data visualization platforms on their computers.
- o Have agency-specific resources for cloud-based compute power that AI developers can harness to train machine learning models with greater speed.

Blueprint for Action: Chapter 6 - Endnotes

¹ Jim Perkins, et al., *Don't Just Copy and Paste: A Better Model for Managing Military Technologists, War on the Rocks* (Aug. 24, 2020), <https://warontherocks.com/2020/08/dont-just-copy-and-paste-a-better-model-for-managing-military-technologists/>.

² These fields were selected from a combination of NSCAI's Third Quarter recommendations and Partnership for Public Service's *Tech Talent for 21st Century Government*. See *Interim Report and Third Quarter Recommendations*, NSCAI (October 2020), <https://www.nscai.gov/previous-reports/Tech-Talent-for-21st-Century-Government>, Partnership for Public Service: *Tech Talent Project* (April 2020), <https://ourpublicservice.org/wp-content/uploads/2020/04/Tech-Talent-for-21st-Century-Government.pdf>.

³ A special government employee is "an officer or employee of the executive or legislative branch of the United States Government . . . who is retained, designated, appointed, or employed to perform, with or without compensation, for not to exceed one hundred and thirty days during any period of three hundred and sixty-five consecutive days." 18 U.S.C. § 202.

⁴ Members of the military reserves typically serve two to three days a month, and one 14-day obligation a year, averaging around 38 days a year.

⁵ Organizations that employ full-time technical experts in temporary positions, such as the United States Digital Service or Defense Digital Service, already exist, and have proven successful. The NRDC is an alternative for experts that cannot or do not want to pursue a full-time route.

⁶ Some administrative functions, such as background checks, security clearance processing, processing tax paperwork, and others, would place an unnecessary burden on local nodes and should be addressed by a central body such as OMB.

⁷ *Uniformed Services Employment and Reemployment Rights Act of 1994*, U.S. Department of Justice (Aug. 6, 2015), <https://www.justice.gov/crt-military/userra-statute>.

⁸ Frank Whitney, *Employment Rights of the National Guard & Reserve*, U.S. Department of Justice (last accessed Jan. 1, 2021), <https://www.justice.gov/sites/default/files/usao-ednc/legacy/2011/04/29/EmploymentRights.pdf>.

⁹ All reservists would apply for security clearances, but this should not imply that reservists would work primarily on classified materials. A large part of the work needed to modernize the government is unclassified.

¹⁰ The Defense Civil Training Corps was created by the National Defense Authorization Act for Fiscal Year 2020. See Pub. Law 116-92, sec. 860, National Defense Authorization Act for Fiscal Year 2020, 116th Congress (2019).

¹¹ Engagement with government officials on Aug. 22, 2019, Feb. 7, 2020, and March 9, 2020.

¹² *CyberCorps: Scholarship for Service*, U.S. Office of Personnel Management (last accessed Jan. 1, 2021), <https://www.sfs.opm.gov/>.

¹³ Code.org (last accessed Jan. 11, 2021), <https://code.org/promote>. See also Oren Etzioni, *What Trump's Executive Order on AI Is Missing: America Needs a Special Visa Program Aimed at Attracting More AI Experts and Specialists*, *Wired* (Feb. 13, 2019), <https://www.wired.com/story/what-trumps-executive-order-on-ai-is-missing/>.

¹⁴ The USDSA should also have the authority to accept gifts, particularly to help fund its establishment.

¹⁵ The Council on Foreign Relations report, *Innovation and National Security: Keeping Our Edge*, recommends creating a digital military service academy. James Manyika & William McRaven, *Innovation and National Security: Keeping Our Edge*, Council on Foreign Relations (September 2019), <https://www.cfr.org/report/keeping-our-edge/>. Our recommendation is for a civilian digital service academy that would not produce any uniformed military personnel.

¹⁶ The five academies include the United States Military Academy, the United States Naval Academy, the United States Coast Guard Academy, the United States Merchant Marine Academy, and the United States Air Force Academy.

¹⁷ Joseph Moreno & Robert Scales, *The Military Academies Strike Back*, The Chronicle of Higher Education (Nov. 12, 2012), <https://www.chronicle.com/article/the-military-academies-strike-back/>. As an example, 5 Secretaries of the Navy, 29 Chiefs of Naval Operations, and nine Commandants of the Marine Corps graduated from the United States Naval Academy.

¹⁸ Each military service academy has a maximum and minimum number of positions available for every available career field, causing some graduates to receive career fields other than their first choice. Similarly, USDSA graduating classes would have a minimum and maximum number of civilian graduates that join each military department or government agency.

¹⁹ The military service academies are accredited by different regional accreditation organizations recognized by the U.S. Secretary of Education and Council for Higher Education. Their engineering programs are generally accredited by the Accreditation Board for Engineering and Technology, Inc.

²⁰ State approval and accreditation are not the same, but both are required.

²¹ Recruitment will rely on private-sector champions to recruit high-profile adjunct faculty that can serve as beacons that will attract additional faculty and high-quality students.

²² The Computing Accreditation Commission on Colleges of Accreditation Board for Engineering and Technology is a nonprofit, ISO 9001 certified organization that accredits college and university programs in applied and natural science, computing, engineering, and engineering technology.

²³ For comparison, since 2001, C:SFS has had 3,600 graduates, or about 189 graduates per year, according to program officials NSCAI spoke with on March 9, 2020.

²⁴ This analysis is based on the NSCAI staff conducting more than 100 interviews with government officials between May 2019 and May 2020. This feedback has emerged as a common theme in nearly all of NSCAI's workforce discussions. See e.g., NSCAI interviews with government officials (June 7, 2019); NSCAI interviews with government officials (May 17, 2019).

²⁵ NSCAI staff interview with a private-sector company (Sept. 9, 2019); NSCAI staff interview with a private-sector company (Sept. 19, 2019); NSCAI staff interview with a private-sector company (April 24, 2020).

²⁶ NSCAI staff interviews with government and private-sector senior leaders (May 6, 2020).

²⁷ *Workforce Now: Responding to the Digital Readiness Crisis in Today's Military*, Defense Innovation Board at 1-7 (2019), https://media.defense.gov/2019/Oct/31/2002204196/-1/-1/0/WORKFORCE_NOW.PDF.

²⁸ NSCAI staff has interviewed several ORSA personnel performing AI-related tasks. All agreed when asked that a separate career field for artificial intelligence or data science is needed. Existing initiatives make some progress, but do not adequately address the lack of career fields for digital talent.

²⁹ The NSCAI staff conducted more than 100 interviews with government officials between May 2019 and May 2020. This feedback has emerged as a common theme in nearly all of NSCAI's workforce discussions. See e.g., NSCAI interviews with government officials (June 7, 2019); NSCAI interviews with government officials (May 17, 2019).

³⁰ NSCAI's *First Quarter Recommendations* included an addition to the Armed Services Vocational Aptitude Battery to test for computational thinking that would help identify aptitude and a test for coding language proficiency that would help identify skill. *First Quarter Recommendations*, NSCAI at 33-35 (March 2020), <https://www.nscail.gov/previous-reports/>. Both tests will be helpful, but will not meet their full utility without digital career fields. In conversations with NSCAI, numerous government officials continuously identified a lack of digital career fields as a key impediment to talent management. See e.g., NSCAI interviews with government officials (June 7, 2019); NSCAI interviews with government officials (May 17, 2019).

³¹ *2020 Interim Report and Third Quarter Recommendations*, NSCAI at 37-38 (October 2020), <https://www.nscail.gov/previous-reports/>.

³² See Chapter 2 of this report for a detailed description of how DoD would implement this plan.

Chapter 7: Establishing Justified Confidence in AI Systems

Blueprint for Action

A Holistic Framework for Ensuring Justified Confidence in AI Systems.

The U.S. Government should align on a common understanding of critical steps needed to ensure justified confidence in AI systems, including confidence in their responsible development and use. The Commission has outlined such a strategy in the *Key Considerations*. The *Key Considerations* provide a framework for the responsible development and fielding of AI that should be adopted by all agencies critical to national security. The framework includes near-term recommendations and topics that agencies should give priority consideration, practices that should be implemented immediately, and policies that should be defined or updated to reflect new AI considerations.

Based on robust feedback from agencies including Department of Defense (DoD), Intelligence Community (IC), Department of Homeland Security (DHS), Federal Bureau of Investigation (FBI), Department of Energy (DoE), Department of State (DoS), and Department of Health and Human Services (HHS), as well as the GSA AI Community of Practice, the *Key Considerations* also outlines areas needing future work and targeted investment to overcome current challenges. Agencies that have already adopted AI principles noted broad alignment between the *Key Considerations* framework and their AI principles. For instance, the framework's recommended practices help operationalize the AI Principles of the DoD and IC¹ and the Principles for Use of AI in Government.²

The implementation of the *Key Considerations'* recommendations for future action will be important not only for agencies, but also for cooperation across the world on the responsible development and fielding of AI.³ Further, while the Commission's mandate led to a focus on recommendations specific to national security entities in our report, many recommendations we elevate in the *Key Considerations* are relevant to the whole country, including other sectors and industry.

Heads of departments and agencies critical to national security should implement the Key Considerations as a framework for the responsible development and fielding of AI systems. Agencies, at a minimum, include the DoD, IC, FBI, DHS, DoE, DoS, and HHS.

Implementing the *Key Considerations* includes developing policies and processes to adopt the framework’s recommended practices, monitoring their implementation, and continually refining them as best practices evolve. While this framework covers dozens of practices that contribute toward an ideal state of responsible development and fielding, some practices will be more critical than others depending on the stakes and context, and complying with them will require different costs and resources. This Blueprint for Action provides details on the key actions from this framework that all departments and agencies critical to national security can and should take now as a priority, and investments and resources that the government should make available to further responsible AI across all agencies. *These span recommendations for Robust and Reliable AI; Human-AI Interaction and Teaming; Testing and Evaluation, Verification and Validation; Leadership; and Accountability and Governance.*

Recommendations for Robust and Reliable AI

Recommendation

Action for the Office of Science and Technology Policy (National AI Initiative Office):

- **Focus federal research and development (R&D) investments on advancing AI security and robustness, to help agencies better identify and mitigate evolving AI system vulnerabilities.** Confidence in the robustness and reliability of AI systems requires insight into the development process and the operational performance of the system. Insight into the development process is supported by capturing decisions and development artifacts for review; insight into operational performance is supported by runtime instrumentation and monitoring to capture details of execution. In both development and operation, there is a need to invest in R&D for better tools to facilitate the capture of needed processes and data. R&D should also advance interpretability capabilities to better understand if AI systems are operating as intended. And R&D should support better characterization of performance envelopes to enable the gradual rollout and adoption of AI systems. “Robust AI” is included among the priority research areas found in Chapter 11 of this report.

Action for all Departments and Agencies

- **Create an AI Assurance Framework.** All government agencies will need to develop and apply an adversarial machine learning threat framework to address how key AI systems could be attacked and should be defended. An analytical framework can help to categorize threats to government AI systems, and assist analysts with detecting, responding to, and remediating threats and vulnerabilities.⁴ This framework must address supply chain threats to data and models as well as adversarial AI attacks.⁵ The framework will support assurance of data authenticity and data and model integrity. “Create an AI Assurance framework” is included among recommendations found in Chapter 1 of this report.

Action for DoD and the Office of the Director of National Intelligence (ODNI):

- **Create dedicated red teams for adversarial testing.** Such red teams should assume an offensive posture, dedicated to trying to break systems and make them violate rules for appropriate behavior.⁶ Because of the scarcity of required expertise and experience for AI red teams, the DoD and ODNI should consider establishing enterprise-wide communities of AI red teaming and vulnerability testing capabilities that could be applied to multiple AI developments. The Commission supports the aligned recommendation by WestExec Advisors that the DoD and ODNI should consider “standing up a national AI and ML red team as a central hub to test against adversarial attacks, pulling together DoD operators and analysts, AI researchers, T&E, CIA, DIA, NSA, and other IC components, as appropriate. This would be an independent red-teaming organization that would have both the technical and intelligence expertise to mimic realistic adversary attacks in a simulated operational environment.”⁷

Actions for Agencies Critical to National Security:⁸

- **To Meet Baseline Criteria for Robust and Reliable AI —**

- o Upgrade development, procurement, and acquisition strategies to ensure that those accountable for the development, procurement, or acquisition of an AI system (e.g., program managers) adopt the following practices:
 - **Consult an interdisciplinary group of experts to conduct hazard analysis and risk assessments.** These should cover, as relevant to the context: potential disparate impact related to unwanted bias; privacy and civil liberties; international humanitarian law; human rights;⁹ system security against targeted attacks;¹⁰ risks of technology being leaked, stolen, or weaponized by adversaries against the U.S.;¹¹ and steps taken to mitigate identified risks. Agencies should specify in their respective strategies who will consult such a group and who will ultimately make final decisions based on the group’s advice.
 - **Improve documentation practices.** Produce documentation describing the data used for training and testing; model(s); other relevant systems (including connections and dependencies within systems); required maintenance (for datasets and models) technical refresh, and when the system is used in a different operational environment. For data, documentation should include how data were sampled, and their provenance. For synthetic data, documentation should also include details on how the data were generated.¹²
 - **Build overall system architectures to limit the consequences of system failure.** Agencies should build an overall system architecture that monitors component performance and handles errors when anomalies are detected; build AI components to be self-protecting (validating input data) and self-checking (validating data passed to the rest of the system); and include aggressive stress testing. As with all high-consequence software systems, where technically feasible, it is important that high-consequence AI systems have overall system architectures that support robust recovery and repair or fail-fast and fail-over to a reliable degraded mode safe system. There should be clear mechanisms for disengaging and deactivating the system when things go wrong.¹³

Recommendations for Human-AI Interaction and Teaming

Recommendation

Action for Department of Defense:

- **Invest in a sustained, multi-disciplinary initiative to enhance human–AI teaming through the Service Laboratories and DARPA.**
 - o This initiative should focus on maximizing the benefits of human–AI interaction; better measuring human performance and capabilities when working with AI systems; and helping AI systems better understand contextual nuances of a situation. Advances in human–machine teaming will enable human interactions with AI-enabled systems to move from the current model of interaction where the human is the “operator” of the machine, to a future in which humans are able to have a “teammate” relationship with machines. Specific funding should be dedicated to research on how to improve human–machine teaming and interaction when it involves human life–safety or lethal deployment of a system. Additional research is urgently needed which should address the following issues, among others: delegation of authority, observability, predictability, directability, communication, and trust.
 - o R&D investment should also focus on the following:
 - Developing improved human performance assessment, an essential element for AI to understand when and how an appropriate AI intervention should be made.
 - Developing new approaches to humans and AI establishing and maintaining common ground in support of collaboration, particularly cognitive collaboration. This encompasses how a newly established human–AI team scaffolds its mutual understanding and then how it extends it to creatively and collaboratively tackle new challenges.
 - Developing new approaches to trust calibration in human–AI teams. This includes helping people understand when AI is approaching or outside the bounds of its competency envelope, and likewise helping machines understand when people are approaching their limits. The two together will help the human–AI team calibrate trust appropriately and shape their interaction for improved team performance.¹⁴ “Enhanced human–AI interaction and teaming” is included among the priority research areas found in Chapter 11 of this report. This recommendation also maps to the overall DoD R&D funding recommendation in Chapter 3 of this report.

Actions for Agencies Critical to National Security:

- **Meet Baseline Criteria for Effective Human–AI Interaction and Teaming —**
 - o National security departments and agencies should clarify policies on human roles and functions, develop designs that optimize human–machine interaction, and provide ongoing and organization–wide AI training.
 - Develop design methodologies that improve our understanding of human–AI interaction and provide specific guidance and requirements that can be assessed.¹⁵ These methodologies should clearly delineate requirements of potential human–AI teaming alternatives and identify whether a proposed solution is likely to meet those requirements or not.

- Designs should mitigate automation bias (that places unjustified confidence in the results of computation) and unjustified reliance on humans as a failsafe mechanism. They should provide accurate cues to the human operator about the level of confidence the system has in its results/ behaviors.
- Ensure policies provide ethical bounds regarding when and where AI is appropriate within a human-AI team in a given context.
 - Policies should identify what functions humans should perform across the AI life cycle; bound assignments and functions, including autonomous functionality; define when tasks should be handed off between a human and machine based on bounds; and require feedback loops to inform oversight and ensure systems operate as expected.
- Provide ongoing training to help the workforce better interact, collaborate with, and be supported by AI systems—including understanding AI tools.¹⁶ As relevant, employees across departments and agencies, and the DoD in particular, should, at a minimum:
 - Gain familiarity with AI tools (e.g., through everyday interaction), including use of AI systems in realistic situations and provide continual feedback to integrate improvements.¹⁷
 - Receive education that includes fundamentals of AI and data science, including coverage of key descriptors of performance and probabilities.¹⁸
 - Receive training on interpreting performance standards and metrics correctly and making informed decisions based on them.¹⁹
 - Gain an understanding of both the fundamental concepts and the high-level concepts in terms of how the system components interact with each other.²⁰
 - Have training to recognize human cognitive biases so that human operators interacting with machines can recognize where they might be succumbing to such bias.²¹
 - Receive ongoing refresher trainings suited to system operators. Refresher trainings are appropriate when systems are deployed in new settings and unfamiliar scenarios, and when predictive models are revised with additional training data as system performance may shift, introducing behaviors that are unfamiliar to operators.²²

Recommendation

Recommendations for Testing and Evaluation, Verification and Validation

Action for the Department of Defense:

- **DoD should tailor and develop TEVV policies and capabilities to meet the changes needed for AI as AI-enabled systems grow in number, scope, and complexity in the Department.**²³

This should address the following elements:

- **Establish a testing and evaluation, verification and validation (TEVV) framework and culture that integrates testing as a continuous part of requirements specification, development, deployment, training, and maintenance and includes run-time monitoring of operational behavior.**²⁴ An AI testing framework should:
 - o Establish a process for writing testable and verifiable AI requirement specifications that characterize realistic operational performance.²⁵
 - o Provide testing methodologies and metrics that enable evaluation of these requirements—including principles of ethical and responsible AI, trustworthiness, robustness, and adversarial resilience.²⁶
 - o Define requirements for performance reevaluation related to new usage scenarios and environments, and distribution over time.
 - o Encourage incorporation of operational usage workflow and requirements from the defined use case into the testing.
 - o Issue data quality standards to appropriately select the composition of training and testing sets.
 - o Support the use of common modular cognitive architectures within suitable application domains that expose standard interface points for test harnessing—supporting scalability through increased automation along with federated development and testing.
 - o Support a cyclical DevSecOps-based approach, starting on the inside and working outward, with AI components, system integration, human-machine interfaces, and operations (including human-AI and multi-AI interactions).
 - o Remain flexible enough to support diverse missions with changing requirements over time.
- **Extend existing and develop new TEVV methods and tools for dealing with complex, stochastic, and non-stationary systems, including the design of experiments, real-time monitoring of states and behaviors, and the analysis of results.** These methods/tools need to account for human-system interactions (HSI) and their impact on system behavior, system-system interactions and their effect on emergent behavior across a group of systems, and adversarial attacks, via both conventional cyberattacks, and nascent perceptual adversarial AI attacks. Risk assurance concepts should be extended beyond simple “stop-light” charts of consequence and likelihood for a risk being realized and leverage tools that support developing assurance cases that present verifiable claims about system behavior and provide reviewable arguments and evidence to support the claims.²⁷
- **Make TEVV tools and capabilities readily available across the DoD,** including downloadable and configurable AI TEVV software stacks.²⁸ In addition, the DoD should ensure tools that support TEVV and reliability and robustness goals are available department-wide including tools for bias detection, explainability, and documentation across the product life cycle (e.g., of data inputs and system outputs).
- **Update existing and create new live, virtual, and constructive test ranges for AI-enabled systems (blending modeling and simulation, augmented reality, and cyber physical system environments).** Upgraded test ranges should include live-virtual-constructive environments, the ability to capture data from testing, and the ability to evaluate data from operations. They should support: 1) The full exploration of potential system states and behaviors over a range of runtimes and fidelity levels;

2) the co-development of AI-system functionality and concepts of operations (CONOPS) associated with human-system and system-system teaming; and 3) a fuller understanding of the impact of adversarial activities undertaken to counter these systems. Build these capabilities upon extensive modeling and simulation (M&S) facilities, human and constructive adversarial “red teams,” virtual and augmented reality enablers, full instrumentation, and post-run big data analytics capability.

- **Support the T&E community by restructuring the processes that underlie requirements specification, system design, T&E itself, and CONOPS development.** This includes continuing DoD investments and policies supporting architecting software-intensive systems using common frameworks and composable subsystems,²⁹ the inclusion of runtime instrumentation (adding the capture of internal states of the system, analogous to a flight data recorder on aircraft) in system design and monitoring during operation,³⁰ the proper curation and protection of data used in training these systems, and a heavy investment in successively sophisticated M&S, starting at the requirements stage and proceeding through development, TEVV, and operator training.

Action for the National Institutes of Standards and Technology (NIST):

- **NIST should provide and regularly refresh a set of standards, performance metrics, and tools for qualified confidence in AI models, data, and training environments, and predicted outcomes.**³¹ Over time, as the science of how to test systems across responsible AI attributes evolves, NIST should provide guidance on:
 - o Metrics to assess system performance per responsible AI attributes (e.g., fairness, interpretability, reliability, robustness) and according to application/context profiles. This should include:
 - Definitions, taxonomy, and metrics needed to enable agencies to better assess AI performance and vulnerabilities.
 - Metrics and benchmarks to assess reliability of model explanations.³²
 - o For each of the metrics and technical measures created, NIST should also provide measurable outcomes against which success can be determined.³³
- **In the near term, NIST should also provide guidance on:**
 - o Standards for testing intentional and unintentional failure modes
 - o Exemplar data sets for benchmarking and evaluation, including robustness testing and red teaming
 - o Defining characteristics of AI data quality and training environment fidelity (to support adequate performance and governance)

In conducting the above, NIST should publish quarterly updates to inform departments and agencies about the trustworthy frameworks, standards, and metrics work it is planning.³⁴

Action for the Office of Science and Technology Policy - National AI Initiative Office:

- **The federal government should increase R&D investment to improve our understanding of how to conduct TEVV.** This is needed to better understand how to efficiently and effectively test AI systems to provide objective assurance to support a justified level of confidence, build checks and balances in systems, and how to monitor and mitigate unexpected behavior in a composed system-of-systems or when systems interact. Such R&D should advance our understanding of how to test system performance across responsible AI attributes (e.g., fairness, interpretability, reliability, and robustness). This recommendation is echoed by the priority research areas found in Chapter 11 of this report, including “TEVV of AI Systems” and “standard methods and metrics for evaluating degrees of auditability, traceability, interpretability, explainability, and reliability.” For more information, see also Chapter 3 of this report.

Actions for Agencies Critical to National Security:

- **To ensure optimal performance of AI systems, national security departments and agencies should:**
 - o Plan for and execute aggressive stress testing of AI components to evaluate error handling and robustness against unintentional and intentional threats under conditions of intended use.
 - o Include testing for blind spots and fairness throughout development and deployment. Testing and validation should be done iteratively at strategic intervention points, especially for new deployments.
 - o Clearly document system performance requirements (including identified system hazards), metrics used for TEVV, deliberations on the appropriate fairness metrics to use, and the representativeness of the test data for the anticipated operational environment.
 - o Conduct red teaming to rigorously challenge AI systems, exploring their risks, limitations, and vulnerabilities including intentional and unintentional failure modes.

Recommendations for Leadership

Recommendation

Actions for DoD, IC, FBI, DHS, DoE, DoS, and HHS:

- **Every department and agency critical to national security and each branch of the armed services, at a minimum, should have a dedicated, full-time Responsible AI Lead who is part of the senior leadership team. Responsible AI Leads must have dedicated staff, resources, and authority to succeed in their roles. Every lead should have at least two full-time staff to effectively fulfill the following:**
 - o The Responsible AI Lead in each department should oversee the implementation of the *Key Considerations* recommended practices alongside the department/ agency’s respective AI principles.³⁵ This includes driving policy development and training programs for the department and internally coordinating Responsible AI leads in the department’s supporting branches or agencies (as applicable) to ensure synergistic implementation of such policies and programs. The department lead should determine the Responsible AI governance structure to ensure centralized and consistent policies³⁶ are applied across the department.

- o The department Responsible AI Lead and those supporting Responsible AI leads should collectively:
 - provide Responsible AI training to relevant personnel;
 - serve as subject matter experts regarding existing and proposed Responsible AI policy and best practices;
 - shape procurement policy and guidance for product managers to ensure alignment with recommended practices and adopted AI principles;
 - build a central repository of Responsible AI work going on in the department, and lessons learned from practical implementation across the department, to help streamline department efforts;
 - ensure interagency knowledge sharing for responsible AI, including iterative sharing of best practices, resources and tools, evolving risks and vulnerabilities, and other lessons learned from practical implementation;
 - annually produce a report for Congress on department resources received, any additional resources needed, and an update on required policy work and implementation of recommended practices.
- o Where possible, centralized assessments and shared learnings should be communicated across a department's elements or branches, to avoid units spending unnecessary and duplicative resources and to accelerate practices that reduce friction in workflows. Responsible AI Leads in each department should consider the Learning, Knowledge, and Information Exchange (LKIE) framework as a way to accelerate organizational knowledge within their department given the need to leverage collective insights that are gleaned from on-the-ground experience where the *Key Considerations* will be put into practice rather than letting the insights sit in silos.³⁷ Furthermore, having Responsible AI "champions"³⁸ who "socialize" this knowledge can help to transfer the knowledge within and across different U.S. Government agencies and components.³⁹
- o Borrowing from the world of cybersecurity, the Lead also should consider coordinating the adoption of an empirically driven prioritization matrix for risk management.⁴⁰

Action for the National AI Initiative Office:

- **In addition to the National AI Initiative responsibilities defined in the National Defense Authorization Act for Fiscal Year 2021 (FY2021 NDAA),⁴¹ the Office should create a standing body of multi-disciplinary experts who can be voluntarily called upon by agencies as a resource to provide advice on Responsible AI issues.** The group should include people with expertise at the intersection of AI and other fields such as ethics, law, policy, economics, cognitive science, and technology including adversarial AI techniques. As the government upskills and diversifies its workforce with AI expertise, this standing body of experts should help fill gaps in multi-disciplinary expertise that can be called upon by agencies as needed for processes including multi-disciplinary risk assessment, human-AI teaming assessments, and red-teaming.
- Leveraging this in-house expertise, and serving as the central resource for best practice sharing across agencies, it should also:

- o Maintain a Learning, Knowledge, and Information Exchange (LKIE) repository to benefit all agencies:
 - A repository compiling insights across agencies (e.g., per the LKIE framework mentioned above) would accelerate organizational knowledge and support the goal of interagency sharing of insights gleaned from on-the-ground practice—rather than letting such insights sit in silos.⁴² These collective insights would be generalized from bright spots of successful AI adoption and from lessons learned from AI adoptions that faced problems in development or use.⁴³ Centralized insights will also provide a resource to help agencies address critical questions that will arise as AI capabilities evolve. Examples of potential critical questions include how to support redress with updated policies and procedures; how to efficiently monitor behavior in operation; and how to effectively measure and address changes introduced by technical refresh. With technical refresh, it is necessary to analyze results carefully. Even if overall performance may be steady or improve after a refresh, the aggregate performance can mask certain parts of the performance envelope where results are significantly skewed and problematic.

Action for Congress:

- **To enable departments and agencies critical to national security to execute Responsible AI work department-wide, and to encourage necessary appointments of Responsible AI personnel, Congress should appropriate an estimated \$21.5 million each fiscal year to fund billets.**

- o Organizations that have high mission complexity and diverse components may need more support staff and/or Responsible AI Leads to be allocated across the organization. The Commission recommends that, at a minimum, the following is needed:
 - For the DoD, a department-wide Responsible AI (RAI) Lead and supporting RAI Leads for each branch of the armed services, with each lead supported by two staff members;
 - For the Intelligence Community, an ODNI RAI Lead and supporting RAI Leads for each IC agency, with each lead supported by two staff members;
 - For the DOE, a RAI Lead and a supporting RAI Lead for the National Laboratories, with each lead supported by two staff members; and
 - For the FBI, DHS, and HHS, a RAI Lead in each respective organization who is supported by two staff members.⁴⁴

Recommendations for Accountability and Governance

Recommendation

Actions for Agencies Critical to National Security:

- **Adapt and extend existing policies to ensure accountability is established and documented across the AI life cycle for any given AI system and its components.⁴⁵**
- **Establish clear requirements about information that should be captured about the development process⁴⁶ (via traceability) and about system performance and**

behavior in operation (via runtime monitoring) to support reliability and robustness as well as auditing for oversight. Instrumentation to support monitoring can contribute to insights about system performance, but must be provided thoughtfully to prevent new openings for external espionage or tampering with AI systems.⁴⁷

- o Guidance should include technical audit trail requirements per mission needs for high-stakes systems.

- **Institute comprehensive oversight and enforcement practices.**

- o Agencies should identify or establish new policies, due to the novelty and advancement of AI technologies, that:
 - allow individuals to raise concerns about irresponsible AI development (e.g., through an ombudsman); and
 - provide layers of human oversight or redundancy so that high-stakes decisions do not rely entirely on determinations made by the AI system.⁴⁸
- o Adapt and extend oversight practices to include reporting requirements⁴⁹ for AI systems; a mechanism to allow for thorough review of the most sensitive and high-risk AI systems (to ensure auditability and compliance with deployment requirements); an appealable process for those found at fault of developing or using AI irresponsibly; and grievance processes for those affected by the actions of AI systems.⁵⁰
- o Establish selection criteria that indicate if and when specific recommended practices (as found in the *Key Considerations*) need to be used according to system and mission risks.
- o Define triggers that would require escalated review of an AI system.

Blueprint for Action: Chapter 7 - Endnotes

¹ See *Key Considerations for Responsible Development & Fielding of Artificial Intelligence Supporting Visuals*, NSCAI (July 2020), <https://www.nscai.gov/wp-content/uploads/2021/01/Key-Considerations-Supporting-Visuals.pdf>.

² See Donald J. Trump, *Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*, The White House (Dec. 3, 2020), <https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-promoting-use-trustworthy-artificial-intelligence-federal-government/>. The Principles for Use of AI in Government do not apply to national security agencies; however, they do apply to agencies the Commission considers critical for national security (e.g., Department of State and Department of Health and Human Services).

³ See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on international cooperation, see the Commission's recommendation for future action in the sections on "Aligning Systems and Uses with American Values and the Rule of Law" and "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁴ There are various public and private efforts ongoing. See for instance the MITRE-Microsoft adversarial ML framework, Ram Shankar Siva Kumar & Ann Johnson, *Cyberattacks Against Machine Learning Systems Are More Common than You Think*, Microsoft Security (Oct. 22, 2020), <https://www.microsoft.com/security/blog/2020/10/22/cyberattacks-against-machine-learning-systems-are-more-common-than-you-think/>; *Adversarial AI Threat Matrix: Case Studies*, MITRE (last accessed Jan. 10, 2021), <https://github.com/mitre/advmlthreatmatrix/blob/master/pages/case-studies-page.md>.

⁵ *NISTIR 8269 (Draft): A Taxonomy and Terminology of Adversarial Machine Learning*, National Institute of Standards of Technology (October 2019), <https://csrc.nist.gov/publications/detail/nistir/8269/draft>.

⁶ See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for red teaming, see the section on "Engineering Practices" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁷ See Michele Flournoy, et al., *Building Trust Through Testing* (October 2020), <https://cset.georgetown.edu/wp-content/uploads/Building-Trust-Through-Testing.pdf>.

⁸ As noted above, the Commission considers these, at a minimum, to include the DoD, IC, DHS, FBI, DoE, Department of State, and HHS.

⁹ For more on the importance of human rights impact assessments of AI systems, see *Report of the Special Rapporteur to the General Assembly on AI and its impact on freedom of opinion and expression*, UN Human Rights Office of the High Commissioner (2018), <https://www.ohchr.org/EN/Issues/FreedomOpinion/Pages/ReportGA73.aspx>. For an example of a human rights risk assessment for AI in categories such as nondiscrimination and equality, political participation, privacy, and freedom of expression, see Mark Latonero, *Governing Artificial Intelligence: Upholding Human Rights & Dignity*, Data Society (October 2018), <https://datasociety.net/wp-content/uploads/2018/10/DataSociety-Governing-Artificial-Intelligence-Upholding-Human-Rights.pdf>.

¹⁰ These can include reidentification attacks. Departments and agencies should use privacy protections such as robust anonymization that can withstand sophisticated reidentification attacks, and when possible, privacy-preserving technology such as differential privacy, federated learning, and ML with encryption of data and models.

¹¹ For exemplary risk assessment questions that IARPA has used, see Richard Danzig, *Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority*, Center for a New American Security at 22 (June 28, 2018), <https://s3.amazonaws.com/files.cnas.org/documents/CNASReport-Technology-Roulette-DoSproof2v2.pdf?mtime=20180628072101>.

¹² Such documentation should support assurances of the authenticity, integrity, and provenance of data.

¹³ See *Making Responsible AI the Norm Rather than the Exception*, Montreal AI Ethics Institute at 9 (Jan. 13, 2021), <https://arxiv.org/pdf/2101.11832.pdf> [hereinafter MAIEI Report] (This includes “building fail safes and backup modes that don’t have to rely on continuous access to the ‘intelligent’ elements and have graceful failures that minimize harm.”).

¹⁴ See Brian Wilder, et al., *Learning to Complement Humans*, Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (2020), <https://www.ijcai.org/Proceedings/2020/0212.pdf>.

¹⁵ For an example of applicable guidelines, see Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI ’19: Proceedings of the CHI Conference on Human Factors in Computing Systems (May 2019), <https://dl.acm.org/doi/10.1145/3290605.3300233>.

¹⁶ For more on training, see the Appendix of this report containing the abridged version of NSCAI’s *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission’s recommendation for training, see the section on “Human-AI Interaction and Teaming” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁷ Such everyday interaction and continual feedback loops will further enhance TEVV.

¹⁸ See the Appendix of this report containing the abridged version of NSCAI’s *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission’s recommendation for training, see the section on “Human-AI Interaction and Teaming” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission). See also MAIEI Report at 7.

¹⁹ MAIEI Report at 7.

²⁰ Id.

²¹ Id.

²² See the Appendix of this report containing the abridged version of NSCAI’s *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission’s recommendation for training, see the section on “Human-AI Interaction and Teaming” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

²³ To the greatest extent possible, DoD should develop TEVV policies and capabilities in coordination with the Office of the Director of National Security.

²⁴ To achieve this, heavy investment is needed that supports requirements generation/traceability, the integration of heterogeneous test data at all stages of testing, and the use of extensive M&S, test automation, and data analytics wherever feasible.

²⁵ This should be framed broadly, providing left/right limits that provide guidance but do not limit innovation.

²⁶ These testing methodologies and metrics should support robust red teaming, meeting the DoD’s particular needs for solutions hardened to adversarial actions.

²⁷ Miles Brundage, et al., *Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims*, arXiv (April 20, 2020), <https://arxiv.org/abs/2004.07213>.

²⁸ TEVV tools and software stacks should be shared across the Department using the AI Digital Ecosystem described in Chapter 2 of this report.

²⁹ Such frameworks for composing testable AI systems should be established and accessed through the AI Digital Ecosystem described in Chapter 2 of this report.

Blueprint for Action: Chapter 7 - Endnotes

³⁰ See e.g., Software Acquisition Pathway Interim Policy and Procedures, Memorandum from the Under Secretary of Defense, to Joint Chiefs of Staff and Department of Defense Staff (Jan. 3, 2020), [https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20\(Software\).pdf](https://www.acq.osd.mil/ae/assets/docs/USA002825-19%20Signed%20Memo%20(Software).pdf) (stating that program managers are required to “achieve ... continuous runtime monitoring of operational software”).

³¹ This recommendation is in line with Congress’ expansion of NIST’s mission regarding AI standards in the FY2021 NDAA, section 5301 to include: “advance collaborative frameworks, standards, guidelines” for AI, “support the development of a risk-mitigation framework” for AI systems, and “support the development of technical standards and guidelines” to promote trustworthy AI systems.” Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

³² “Documentation of the assumptions and limitations of the benchmarks so created will also be essential in helping those utilizing them to make sure they will get the intended intelligence from it rather than becoming falsely confident about the system.” MAIEI Report at 9.

³³ MAIEI Report at 9.

³⁴ Doing so will enable departments and agencies to plan and prioritize any internal standards work accordingly (e.g., avoiding redundant or obsolete efforts).

³⁵ For each of the metrics and technical measures mentioned in the Key Considerations, it will be important to have measurable outcomes against which success can be determined. See MAIEI Report at 9.

³⁶ This includes, for example, “Accountability and Governance” policy work identified below in this Blueprint for Action.

³⁷ MAIEI Report at 11-16.

³⁸ “AI champions” are a cross-functional group of ambassadors, who can, for example, consider ways to operationalize AI ethical principles and serve as internal advocates and evangelists for responsible AI. See *Department of Defense Joint Artificial Intelligence Center Responsible AI Champions Pilot*, DoD (last accessed Feb. 3, 2021), https://www.ai.mil/docs/08_21_20_responsible_ai_champions_pilot.pdf; Tim O’Brien, et al., *How Global Tech Companies can Champion Ethical AI*, World Economic Forum (Jan. 14, 2020), <https://www.weforum.org/agenda/2020/01/tech-companies-ethics-responsible-ai-microsoft/>.

³⁹ MAIEI Report at 12.

⁴⁰ MAIEI Report at 20-23.

⁴¹ Pub. L. 116-283, Div. E., Title LI, sec. 5102, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁴² MAIEI Report at 11-16.

⁴³ For instance, this could include communication of failure modes (e.g., when a system produces a formally correct, but unsafe outcome), and instances to establish a shared understanding of how and where the systems go wrong. Leveraging this, agencies should tap into USG network-wide expertise to address those failures. See Ram Shankar Siva Kumar, et al., *Failure Modes in Machine Learning*, Microsoft (Nov. 11, 2019), <https://docs.microsoft.com/en-us/security/engineering/failure-modes-in-machine-learning>; MAIEI Report at 7.

⁴⁴ Collectively, considering both Responsible AI Leads and supporting staff, this recommendation proposes 21 full-time employees (FTEs) for the DoD; 54 for the IC; 3 for the FBI; 3 for DHS; 6 for DoE; 3 for HHS; and 3 for DoS.

⁴⁵ As noted in the *Key Considerations*, agencies should determine and document who is accountable for a specific AI system or any given part of an AI system and the processes involved with it. This should identify who is responsible for the development or procurement; operation (including the system's inferences, recommendations, and actions during usage) and maintenance of an AI system; as well as the authorization of a system and enforcement of policies for use. See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for accountability, see the section on "Accountability and Governance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁴⁶ For a list of recommended information that documentation should note about system development, see the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendations for traceability, see the *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁴⁷ For example, "APIs are 'doors' to access digital infrastructures thus, the security and resilience of digital environments will also depend on the robustness of the API infrastructure." V. Lorenzino, et al., *Application Programming Interfaces in Governments: Why, What and How*, European Union Joint Research Centre (2020). <https://ec.europa.eu/jrc/en/publication/eur-scientific-and-technical-research-reports/application-programming-interfaces-governments-why-what-and-how>.

⁴⁸ See Frances Duffy, *Ethical Considerations for Use of Commercial AI*, John Hopkins Applied Physics Laboratory at 31 (Dec. 2020). For example, DoD Directive 3000.09 requires human oversight in the targeting and execution process for lethal autonomous weapons. See *DoD Directive 3000.09: Autonomy in Weapons Systems*, U.S. Department of Defense (May 8, 2017), <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf>.

⁴⁹ For example, reporting risk and impact assessment, steps taken to mitigate such risks, and system performance during testing and fielding.

⁵⁰ As with all consequential software systems, developers and adopters of consequential AI systems must adapt and extend existing support for oversight, audit, reporting, and appealable accountability for developing or using systems irresponsibly, and a redress process where appropriate for those affected by system actions. Existing frameworks must be tailored to reflect issues of concern with AI-based systems (particularly based on machine learning). These issues of concern are discussed in more detail in the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendations for accountability and governance, see the *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission)

Chapter 8: Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security

Blueprint for Action



The U.S. needs an approach for adopting AI domestically for national security that upholds and bolsters respect for democratic values, including privacy, civil liberties, and civil rights. Such an approach must strengthen, provide, and show leadership with regard to: 1) transparency; 2) approaches for AI system development and testing; 3) the ability to contest AI decisions; 4) oversight over AI development and use; and 5) legislative and regulatory controls on how AI is used. Our recommendations include immediate actions that the President, the Congress, and agencies should take; a comprehensive assessment by a Task Force that leads to reforms for AI governance and oversight; and areas for continued work. The recommendations are aimed at assuring that AI systems used by national security agencies uphold democratic values. Secondly, the adoption of these recommendations can earn and inspire public confidence, both domestically and abroad, in uses of AI by national security agencies.

Recommendation

Recommendation Set 1: Increase Public Transparency about AI Use through Improved Reporting

Actions for Congress:

- **For AI systems that involve U.S. persons, require AI Risk Assessment Reports and AI Impact Assessments to assess the privacy, civil liberties and civil rights implications for each new qualifying AI system or significant system refresh.**
 - o The Commission proposes Congress require elements of the Intelligence Community (coordinated by the Office of the Director of National Intelligence (ODNI)) as well as the Department of Homeland Security (DHS) and the Federal Bureau of Investigation (FBI), to prepare and publish an AI Risk Assessment Report and conduct AI Impact Assessments to assess the privacy, civil liberties, and civil rights implications of each new qualifying AI system or significant system refresh. The Commission recognizes the current requirements for privacy impact

assessments and civil liberties impact assessments done at agencies as required by current statute. AI-related technologies may be reviewed by these, but are not fully/adequately captured by the current assessments. The Commission's recommendation intends to augment these requirements.

- o The AI Risk Assessment Report and AI Impact Assessment would be required for “new qualifying AI systems” and for “significant system refreshes.” The Commission proposes that the Task Force described later in this Blueprint be charged with determining the decision procedures for identifying which AI systems and significant system refreshes would require AI Risk and Impact Assessment Reports.
- o The intent of the AI Risk Assessment Report and AI Impact Assessment is to ensure potential impacts are considered and mitigated while avoiding an unnecessary increase in compliance burdens.
 - Legislated frameworks for ensuring effective and pragmatic risk mitigation (with the ability to categorize systems per risk and determine the appropriate mitigations if any) exist in other models that can be used as a template (e.g. FISMA).
- o The *AI Risk Assessment Report* should include a detailed analysis of system implications for, and *steps to mitigate and track risks (e.g., through metrics)* to:
 - Freedom of expression (e.g., is the AI-enabled surveillance targeting people because of their First Amendment protected activity or is the AI-enabled government surveillance causing or may potentially cause a chilling effect?);
 - Equal protection (e.g., is the AI-enabled surveillance biased toward a protected class? What are the likely effects the new technology or program will have on key demographics such as race, gender, or disability?);
 - Privacy (e.g., is a warrant required for the government action? Are minimization and query processes sufficient/satisfactory?);
 - Redress and due process (e.g., what mechanisms exist, or limitations have been accepted, for providing redress for adverse government actions taken based on information generated by the AI system?); and
 - The assessment should account for the environment in which the AI system will be deployed, including its interactions with other AI tools and programs that collect personally identifiable information (PII).
- o *AI Impact Assessment* should be made available periodically, but no less than annually, to the agency's Privacy and Civil Liberties (PCL) Office to determine the degree to which a qualifying AI system remains compliant with the constraints and metrics established in the Risk Assessment. AI Impact Assessments should be based on outcomes, impacts, and metrics collected during system use, and determine if the existing validation processes should be improved.
- o *Resources and staffing.* PCL Offices should assess the resources, including staff, needed to adequately complete the above. Agency heads should support additional resourcing for PCL Offices as part of the annual budget process.
- o *Disclosure notices.* Congress should require ODNI, DHS, and the FBI to review non-public and/or classified AI programs once the program is shut down for declassification and/or disclosure.

Action for the President:

- **Should Congress not require new privacy, civil liberties, and civil rights reporting (as identified above), conduct AI Risk Assessment and AI Impact Assessment Reports as described above.**

Actions for DHS and the FBI:

- **DHS and the FBI should impose new obligations for System of Record Notices (SORNs) and Privacy Impact Assessments (PIAs) specific to AI systems to ensure that they provide richer information.**
 - o SORNs and PIAs should provide a holistic picture about the collection, use, and storage of personal information by any AI system, including its connections to existing systems and accounting for the layering of different surveillance technologies where applicable. Agency practices do not sufficiently support the production of SORNs and PIAs that adequately depict how AI systems collect, use, and store personal information.¹
 - o DHS and the FBI should require that all PIAs include description of the algorithm(s) used and purpose of the algorithm(s); the potential for inferring additional information about individuals from the aggregation of multiple data sources; and importantly, the measures that will be used to address these risks.
 - o DHS and the FBI should require that SORNs provide more specificity in describing types of data collected, data sources and the connections between data sources, and who will use such data and why.
- **DHS and the FBI should take steps to increase public transparency about the AI systems they employ.**
 - o DHS has recently started an effort to improve transparency, and those efforts should be prioritized and assessed as they are implemented.²
 - o The FBI should implement similar reforms to improve transparency.
- **DHS and the FBI should make their websites easier for the public to navigate and ensure the websites are regularly updated.** Privacy, Civil Liberties, and Civil Rights Risk and Impact Assessment Reports, related semiannual reports, PIAs, and SORNs should be located in a central place; have clearly marked dates next to the title, and chronologically ordered, and published in a timely manner. DHS and the FBI should seek public comments annually about the navigability of their websites and potential improvements.

Recommendation

Recommendation Set 2: Develop & Test Systems per Goals of Privacy Preservation and Fairness

Actions for the President:

- **Through Executive Order, the President should require the Director of National Intelligence, the Secretary of Homeland Security, and the Director of the FBI to take the following actions:**

- **Implement steps to mitigate privacy, civil liberties, and civil rights risks associated with any AI system on an iterative basis and require documentation of all accepted risks.**

- o In implementing steps to achieve this objective, the Commission recommends that ODNI, DHS and the FBI adopt practices from the Key Considerations. In particular:
 - Use privacy protections such as robust anonymization that can withstand sophisticated reidentification attacks, and when possible, privacy-preserving technology such as differential privacy, federated learning, and machine learning (ML) with encryption of data and models.³
 - Mitigate bias in development and testing. For development, conduct stakeholder engagement to establish consensus on the definition of fairness metrics to be used for the specific development and identify necessary constraints on system behavior to protect civil rights and avoid inequitable outcomes.⁴ In testing, confirm that identified constraints are enforced.⁵ Testing to expose unintended bias should include testing for and documentation of different types of error rates (e.g., differences in false positive or false negative rates) or disparate outcomes across demographics.⁶
 - Use AI-tools to support assessing fairness (e.g., industry tools cited in the *Key Considerations*).⁷
 - Ensure the MLOps toolchains include routine calibration of agreed-upon fairness metrics throughout continuous development and integration.⁸
 - Assess model performance and system impact during fielding on an ongoing basis, including emergent behavior, to ensure compliance with privacy, civil rights, and civil liberties objectives.⁹

- **Designate an office, committee, or team in each agency to conduct a pre-deployment review of AI technologies that will impact privacy, civil liberties, and civil rights, including relevant documentation.**

- o This should include review in advance of their deployment and for compliance over the life span of the system.¹⁰ An office in each Intelligence Community agency, DHS, and the FBI should be equipped to assess data, model, and system documentation, and testing results of technologies per their intended use.
- o In undertaking this review, the Commission recommends the designated office use the *Key Considerations*.

Actions for Congress:

- **Establish third-party testing center(s) to allow independent, third-party testing of national-security-related AI systems that could impact U.S. persons.**

- o Congress should fund NIST to create a Third-Party AI Testing Lab program under the NIST National Voluntary Laboratory Accreditation Program.¹¹

- o The third-party test mechanism's mandate should be to:
 - Tailor metric assessment per agency mission and authorities;
 - Develop an approach for testing both software products that can be installed in a test facility and cloud-based services;
 - Establish binding data dissemination agreements with stakeholders of the system to be tested (e.g., the agency requesting testing and relevant vendors and data owners);
 - Collaborate with the agency seeking testing to reach consensus on how to handle the test data provided and the test results and analyses.¹²
- o Third-party test center(s) should allow government vendors to share proprietary data without fear of it being exposed to competitors; and offer the benefits of an aggregated view of performance across a sector or collection of corporations and aggregated best practices.
- o Third-party test center(s) should be used by agencies prior to procuring or fielding high-consequence systems that impact U.S. persons, and use should be considered to overcome in-house testing limitations.
- **Require the Department of Justice (DOJ), in consultation with the Privacy and Civil Liberties Oversight Board (PCLOB), to develop binding guidance for the use of third-party testing (e.g., thresholds for high-consequence systems or unprecedented factors) of AI systems.¹³**
 - o This should include criteria for when an AI system may pose high enough risk for privacy, civil liberties, and civil rights that it would trigger a testing requirement by a third party. In forming such guidance, PCLOB and the DOJ should consult with PCL Officers in federal agencies.

Acknowledgment of continued work for the technical community and legal experts.

There are significant unresolved tensions between various technical approaches to preserving civil rights and civil liberties and current and anticipated legal frameworks. For example, scholars have expressed concern “that technical and legal approaches to mitigating bias will diverge so much that laws prohibiting algorithmic bias will fail in practice to weed out biased algorithms and technical methods designed to address algorithmic bias will be deemed illegally discriminatory.”¹⁴ Continued work in the technical, legal, and policy domains is required to find a consensus balance that addresses technical approaches to preserving privacy, civil liberties, and civil rights and evolving policy.

Recommendation

Recommendation Set 3: Strengthen the ability of those aggrieved by AI to seek redress and have due process.

Actions for FBI and DHS:

- **The FBI and DHS should each conduct a review of its respective policies and practices related to AI technology to ensure that parties aggrieved by government**

action involving the use of AI, including through system actions or misuse, can seek redress and clearly know how to do so. At least annually, the FBI and DHS shall assess if updates or changes are required to their respective reviews.

- o This review should determine whether notice of AI use in decision-making is adequately provided to aggrieved parties to enable redress, as well as the degree of auditability and interpretability needed to contest.
- o The FBI and/or DHS review team—which must include the Offices of Privacy and Civil Liberties—should submit recommendations to their respective agency heads for any regulatory and/or policy changes necessary to amend existing redress mechanisms to reflect issues raised by the use of an AI-enabled system.
- o The Attorney General, working with the Director of the FBI, and the Secretary of Homeland Security, respectively, should direct appropriate actions to ensure that each agency:
 - provides adequate redress, based on the recommendations of the review; and
 - provides the public with clear, updated guidance on how to seek redress in situations covered by the review, including by posting relevant information on their websites.

Actions for the Attorney General:

- **Issue federal guidance on AI and due process. This guidance should describe how relevant agencies should safeguard the due process rights of U.S. persons when AI use may lead to a deprivation of life or liberty.** This should include what obligations agencies have to disclose on its use of AI¹⁵ to a criminal defendant in a timely manner prior to trial or hearing (i.e., notice obligations), including the role that AI played leading to an arrest, charge, or criminal prosecution. Such guidance should be incorporated into agency operational guidelines.

Acknowledgment of continued work by the judicial and/or legislative branches:

The above actions should ensure that agencies receive clear guidance on AI-related redress and due process¹⁶ in the interim as Congress and/or the courts weigh in on federal requirements. Continued work will be needed to provide baseline guidance with the evolution of AI capabilities and their application,¹⁷ and to address open questions on the federal rules of evidence and criminal procedure as they relate to AI.¹⁸

Recommendation Set 4: Strengthen Oversight and Governance Mechanisms to Address Current and Evolving Concerns

Recommendation

Actions for Congress:

- **Strengthen the Privacy and Civil Liberties Oversight Board's (PCLOB) ability to provide meaningful oversight and advice to the federal government's use of AI-enabled technologies for counterterrorism purposes.** To achieve this, Congress

should provide for a targeted expansion of PCLOB's authorities and appropriations as follows:

- o *Awareness of AI programs.* As part of PCLOB's authority to access all relevant material from agencies, agencies should be required to provide PCLOB notice prior to the fielding or repurposing of an AI system, as well as any associated privacy, civil liberties, and civil rights impact assessments.
- o *Visibility into technology.* Agencies should be required to provide to PCLOB, upon PCLOB's request, specific information about technology used in any AI system, including: the data used for AI systems (e.g., documentation regarding the data collection processes for AI-enabled tools and programs, including disclosure and consent processes); models used (and supporting model documentation regarding training and testing); and model repurposing (beyond that context for which it was trained/approved).
- o *Resources and other organizational requirements.* PCLOB requires an increase to its resources, both in terms of talent and funding, to achieve its mission and manage its portfolio as AI adoption increases. PCLOB should provide Congress with a self-assessment of its resources and organizational structure given the expected increase of AI-related programs that fall under its current mandate and responsibilities.
- **Empower DHS Offices of Privacy and Civil Rights and Civil Liberties.** Congress should bolster the roles of DHS' Office of Privacy and Office of Civil Rights and Civil Liberties by requiring the Chief Civil Rights and Civil Liberties Officer, in coordination with the Privacy Officer, to play an integral role in the legal and approval processes for the procurement and use of AI-enabled systems, including associated data of machine learning systems in DHS. As part of this legislation, the Privacy and Civil Rights and Civil Liberties offices should report back to Congress concerning additional staffing or funding resources that are required to satisfy this mandate.

Action for the Secretary of Homeland Security:

- **Ensure the Privacy Officer and the CRCL Officer receive permanent seats in the new DHS enterprise-wide AI Coordination and Advisory Council.** Such appointments are needed in order to meaningfully satisfy the DHS AI Strategy objective titled, "Formalize AI Governance Processes at DHS."¹⁹

Actions for the President:

- **Through Executive Order, require stronger coordination and alignment among oversight and audit organizations through creation of an interagency working group focused on oversight and audit.** Voluntary compliance by agencies with AI documentation and testing requirements should be supported by rigorous, technically informed oversight. To achieve this and overcome current auditing impediments, a standing body (e.g., an interagency working group) should align and coordinate to enhance AI oversight and audit with respect to privacy, civil liberties, and civil rights. This includes system auditability such that the government can monitor and trace the steps that produced a system's output,²⁰ and auditing to ensure systems are not being misused.

- o *Composition:* Organizations should include the Department of Justice Intelligence Oversight Section; Office of the Inspector General of the Intelligence Community; the Government Accountability Office; the Privacy & Civil Liberties Oversight Board; Civil Liberties and Privacy Offices of national security agencies; the National Security Council, and the Office of Science & Technology Policy.
- o *Function:* The interagency working group should provide a forum for members to substantively and regularly address and share information. The working group should:
 - Develop an inventory of the types of AI-relevant oversight and audit currently performed by and anticipated by the participant organizations.
 - Develop an inventory of specific capabilities developed in each organization to address AI oversight and audit.
 - Assess available AI-enabled tools that can be adapted to support more effective and efficient oversight and audit.
 - Tools that support financial audit²¹ and model risk management²² are examples of advances in applying AI to improve the efficiency and scalability of audits that should be reviewed for adoption.
 - Identify priority investment requirements for each organization to address current needs.
 - Identify priority research topics for open S&T gaps in supporting AI oversight and audit, including research gaps in applications of AI in support of privacy and civil liberties (e.g., ML techniques for classification, recommendation, anomaly detection, and other applications)²³ and extending tools such as those that support financial audits and model risk management;
 - Recommend policy or legislative changes for specific authorities granted to the individual organizations.
 - Address mission and focus overlap among representative organizations.
 - Issue reports, at a minimum annually, on key oversight and audit activities as well as S&T gaps.

Action for the President or Congress:

- **Establish a task force to assess the privacy and civil rights and civil liberties implications of AI and emerging technologies.**

The goal of the task force would be to identify gaps and make recommendations to ensure that uses of AI and associated data in U.S. government operations comport with U.S. law and values, and to study organizational reforms that would support this goal. Specifically, it should assess existing policy and legal gaps for current AI applications and emerging technologies, and make recommendations for:

- legislative and regulatory reforms on the development and fielding of AI and emerging technologies;²⁴ and

- institutional changes to ensure sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies.

As mentioned in Chapter 8 of this report, the advancement of AI requires a forward-looking approach to oversight that anticipates the continued evolution and adoption of new technologies, and better positions the government to manage their employment responsibly well into the future. The Commission assesses that, to achieve this goal, a new task force is needed.

The Commission recommends that the President or Congress create a task force with the proposed membership, structure, function, and priorities identified below.

For expediency, the President should:

- **Issue an Executive Order that creates a task force charged with recommending reforms for AI governance and oversight.**
 - o *Membership and structure.* The President should create a task force in the Executive Office of the President to develop recommendations on ensuring adequate AI governance and oversight. The President should designate a senior official to lead the task force. Members should include the heads of OMB, NIST, PCLOB, and the GAO; and Chief Civil Liberties and Privacy Officers and Inspectors General of all national security agencies. In addition, the task force should include representatives from civil society (including organizational leaders with expertise in privacy, civil liberties, and civil rights), industry, and academia. The National AI Advisory Committee Subcommittee on AI and Law Enforcement should also be represented.²⁵
 - o *Function.* The task force should be charged with the following responsibilities:
 - Conducting a macro assessment of the privacy and civil rights and civil liberties implications of the capabilities of AI and emerging technologies;
 - Making recommendations for legislative and regulatory reforms on the development and fielding of AI and emerging technologies, including associated data, in the following key areas:
 - *Privacy, Civil Liberties, and Civil Rights (P/CLCR) reporting.* Binding guidance on when the IC, DHS, and FBI should prepare and publish an AI Risk Assessment Report and AI Impact Assessments, specifically what constitutes a qualifying AI system or significant system refresh (as discussed in the first recommendation of Chapter 8 of this report).
 - *Biometric technologies.* This should include baseline standards for federal government use of biometric identification technologies, including but not limited to, facial recognition.
 - o To address the urgent need for baseline standards and safeguards regarding facial recognition, this should include assessing gaps in federal legislation, gathering input from agency stakeholders (and their legal counsel) currently using facial recognition for national security missions; privacy, civil liberties, and civil rights experts inside and outside of government, including PCLOB; and from the public at large in order to make facial recognition legislation recommendations.

- o Beyond facial recognition, guidance will be needed regarding other biometric identification tools including voiceprints.
- *Government procurement of commercial AI products.* This should include contractual best practices for ensuring industry AI products (including associated data) procured by the government uphold privacy, civil liberties, and civil rights expectations (including privacy, information security, fairness/non-discrimination, auditability, and accountability). This should include third-party requirements that should be incorporated into procurement terms that speak to responsible AI objectives, including for testing validation.²⁶ Consideration should be given to both government-off-the-shelf and commercial-off-the-shelf (COTS) procurement.²⁷
- *Data privacy and retention.* Updates to and reforms of government data privacy and retention requirements to address AI implications.
- Making recommendations for institutional changes to ensure sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies.
 - Evolving AI capabilities are poised to challenge existing expectations for privacy, civil liberties, and civil rights.²⁸ In light of this, the task force should assess the utility of a new entity within the federal government to regulate and provide government-wide oversight of AI use by the federal government.
 - In evaluating options for a new entity, the task force should consider the following:
 - o Authorities and resources necessary for the new entity to provide ongoing guidance and baseline standards for:
 - The federal government's development, acquisition, and fielding of AI technologies to ensure they comport with privacy, civil liberties, and civil rights law and values, and to include guardrails for their use and disallowed outcomes²⁹ to be incorporated in policy and embedded in system development; and
 - Transparency to oversight entities and the public regarding the Federal Government's use of AI systems and the performance of those systems.
 - o Existing interagency and intra-agency efforts to address AI oversight; and
 - o The unique needs of national security, law enforcement, and other government missions with respect to AI systems and potential implications for privacy, civil liberties, and civil rights, and civil liberties.
 - After considering the potential utility of a new organization, make recommendations on organizational placement and structure, composition, authorities, and resources needed.
- Assessing ongoing efforts to adapt regulation of the private sector's AI adoption,³⁰ and as appropriate, consider and recommend institutional or organizational changes to facilitate adequate regulation of commercial development and fielding of AI and associated data.

- o *Reporting.* The task force should issue a report to the President with its legislative and regulatory recommendations on a rolling basis, but no later than within 180 days of its creation. It should issue a report to the President with its recommendations for organizational changes within one year of its creation. The Commission recommends that the report be provided to Congress to ensure transparency and assist Congress in examining these critical issues.


- **In the alternative, Congress should mandate the existence of this task force as outlined above.**

Acknowledgment of continued work to update and clarify legal frameworks on key issues in data protection and data privacy:

A comprehensive approach to upholding privacy and civil liberties in the AI era requires tackling several large, unresolved policy and legal questions regarding data protection and data privacy. Detailed recommendations on these issues would extend beyond the scope of this Commission's mandate, but we identify them here in order to urge further study and congressional action.

- *Legal concerns over federal use of third-party data.* Congress and/or the Judiciary should assess the adequacy of current legal constraints over the federal government's obtainment and use of third-party data, including data acquired from data brokers. Either through evolving case law or legislation, agencies would benefit from clarity surrounding the Fourth Amendment's application on third-party data.³¹ In the meantime, agencies should provide transparency on their respective policies and legal basis for accessing and using commercial data.³²
- *National data protection standards.* Data privacy policies and standards that apply to government agencies alone will be inadequate, and in some cases may harm national security interests.³³ This is particularly important considering how adversaries (both state and non-state actors) can access and use data collected about U.S. persons. As Congress considers proposals for national data security and privacy protection, it should ensure any future legislation addresses the issue of microtargeting. As noted in Chapter 1 of this report, AI systems will create new capabilities for state actors to target individuals with precision as well as numerous aspects of our society like cities, supply chains, universities, corporations, infrastructure, and financial transactions. Strong data privacy protections will be necessary for a multitude of reasons, including to shield the United States from this new phenomenon.
- *National framework for use of biometric technologies.* In the absence of federal legislation regulating the use of facial recognition, the existing patchwork of state and local laws and regulations creates a number of difficulties for government officials, industry, and the public. This has led to actions including: companies prohibiting the sale of facial recognition to law enforcement,³⁴ and local government bans on the use of facial recognition have emerged from coast to coast.³⁵ The lack of a consistent federal approach is also a liability for national security agencies when best practices are not used locally.³⁶ In developing regulation, it will be critical that policy and legislation account not only for facial recognition, but other types of biometric identification that, when combined with other AI technology, can introduce additional concerns.³⁷

Unregulated and Legal Data Collection & Brokering for AI-enabled Predictions and Identification



Full Name, Age 35

Female, Born July 4, 1985



Known Data

- Name
- Gender
- Age/Birthdate
- Birthplace
- Relationship Status
- Contacts, Family, & Associates
- Address
- Address History
- Phone Numbers
- Occupation/Employer
- Professional/Business Licenses
- Salary/Wealth Data
- Registered Political Party
- Voting History
- Court Filings
 - Bankruptcy Records
 - Arrest Records/Mugshot
 - Marriage/Divorce Filings
- E-mail Addresses
- Browsing History
- Shopping History
- Driver's License Number
- Accident History
- Education History
- Geolocation History

Data Brokers

- Web Scraping
- Mobile App Data, End-User License Agreements, Collected Data
- Public Records
- Social Media Scraping



Inferred Data

Thousands of data points used to create statistical inferential profile of an individual.

- Alcohol/Tobacco User?
- Sexual Activity?
- Social Media Influencer?
- Influenced by Social Media?
- Government/Military?
- Mental Health Status?

Blueprint for Action: Chapter 8 - Endnotes

¹ For instance, a recent DHS IG report criticizes the DHS Privacy Office for not establishing controls to ensure that privacy compliance documentation is complete and submitted as required, and specifically noted DHS had not performed required periodic reviews for new and evolving privacy risks. DHS IG, *DHS Privacy Office Needs to Improve Oversight of Department-wide Activities, Programs, and Initiatives*, OIG-21-06, (Nov. 4, 2020). Civil society members have noted that PIAs and SORNs are often too opaque to be helpful, and that agencies sometimes try to shoehorn new data collections under older SORNs and PIAs. See *Comments of the Electronic Frontier Foundation Regarding System of Records Notices 09-90-2001, 09-90-2002*, Electronic Frontier Foundation (Aug. 17, 2020), https://www.eff.org/files/2020/08/17/2020-08-17_-_eff_comments_re_hhs_regs_re_covid_data.pdf (criticizing two SORNs issued by the Department of Health and Human Services during the pandemic, as “overly vague in describing the categories of data collected, the data sources, and the proposed routine uses of the data”).

² The Commission acknowledges DHS’ steps to improve public records as noted in the DHS AI Strategy: “Future AI systems implemented by DHS will require a public release of system information with appropriate exceptions for certain sensitive military and intelligence systems, and some exceptions for law enforcement activities. DHS will produce a framework for releasing AI system information and a process for public comment.” See *U.S. Department of Homeland Security Artificial Intelligence Strategy*, U.S. Department of Homeland Security at 14 (Dec. 3, 2020), <https://www.dhs.gov/publication/us-department-homeland-security-artificial-intelligence-strategy>.

³ To support agencies in this goal, federal R&D investment should continue to advance the state of the art for preserving personal privacy. For information regarding the critical AI research areas the Commission recommends OSTP prioritize, see the Chapter 11 Blueprint for Action. Agencies should also assign responsibility for assessing the state of the practice and encouraging federated learning and anonymization pilots for government databases used in machine learning developments (e.g., to Chief Data Officers at each agency).

⁴ Development practices should also include documenting trade-offs made, including optimizations that cause a trade-off in performance across fairness metrics.

⁵ For instance, constraints about proxies for national origin or protected classes used for rules-based system predictions.

⁶ These include: 1) Documenting operating thresholds including those that yield different true positive and false positive rates or different precision and recall across demographics; (2) Assessing the representativeness of data and model for the specific context at hand; (3) Using tools to probe for unwanted bias in data, inferences, and recommendations; (4) Testing for fairness and articulating the approach, performance, and metrics used. For an extensive list of practices, see the Appendix of this report containing the abridged version of NSCAI’s *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission’s recommendations to mitigate bias in development and testing, see the *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁷ Examples of tools available to assist in assessing and mitigating bias in systems relying on machine learning include Aequitas by the University of Chicago, Fairlearn by Microsoft, AI Fairness 360 by IBM, and PAIR and ML-fairness-gym by Google. Microsoft’s AI Fairness checklist provides an example of an industry tool to support fairness assessments. See Michael A. Madaio et al., *Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI*, CHI 2020 (April 25-30, 2020), <http://www.jennwv.com/papers/checklists.pdf>.

⁸ A widely used Industry example of a fairness metric is Equality of Opportunity (EEO), defined in *Machine Learning Glossary: Fairness*, Google Developers (Feb. 11, 2020), <https://developers.google.com/machine-learning/glossary/fairness>. Note that EEO is suited for some contexts and a poor fit for others—this is why careful deliberation of the operational metrics for fairness must be established early in the development process.

⁹ Select practices include: 1) Assessing statistical results for performance over time to detect emergent bias; 2) recurrent testing and validation at strategic milestones, especially for new deployments and classes of tasks; and 3) Continuously monitoring AI system performance, including the use of high-fidelity traces to determine if a system is going outside of acceptable parameters (e.g., for fairness and privacy leakage) pre-deployment and in operation. For an extensive list of practices, see the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for maintenance and deployment, see the section on "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁰ ML systems in particular require ongoing assessments of privacy and fairness assurances, including the specific definition of fairness being assumed.

¹¹ This requires the creation of an AI TEVV handbook, a culmination of applied research, to create the testing protocols that should be carried out by third-party testing lab(s) and the accreditation procedures by which labs can become certified.

¹² In some cases, exposure of test results could reveal weaknesses in a national security system that could be exploited by an adversary.

¹³ As noted in *Ethical Considerations for Commercial Use of AI*, "rigorous testing is particularly important for high-risk applications, and standards should be established to determine the nature of those applications." See Frances Duffy, *Ethical Considerations for Use of Commercial AI*, Johns Hopkins Applied Physics Laboratory (Dec. 2020).

¹⁴ Alice Xiang, *Reconciling Legal and Technical Approaches to Algorithmic Bias*, Tennessee Law Review at 7 (July 13, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3650635. See also Zachary Lipton, et al., *Does Mitigating ML's Impact Disparity Require Treatment Disparity?*, arXiv (Jan. 11, 2019), <https://arxiv.org/abs/1711.07076>. (Some approaches to mitigate disparate outcomes explicitly make use of membership in protected classes such as race or gender, and are demonstrably more equitable than comparable algorithms that are "blind" to protected classes.)

¹⁵ Disclosure requirements should be specific to each application of AI. See Frances Duffy, *Ethical Considerations for Use of Commercial AI*, Johns Hopkins Applied Physics Laboratory at 31 (December 2020). ("Appropriate disclosure requirements should be created for the purposes of traceability in a court case or for the government's own internal use.")

¹⁶ As noted in the *Key Considerations*, existing policies for contestability should be assessed and updated as needed to ensure accountability and to mitigate errors through feedback loops. See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation to adopt policies to strengthen accountability and governance, see the section on "Accountability and Governance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁷ Due process rights require that individuals have the ability to meaningfully challenge a decision made against them. In federal criminal trials, this includes having the government's explanation of how an unfavorable decision was reached, so it can be contested. In cases where AI-assisted or AI-enabled decisions are made, certain AI techniques will be less conducive to due process. See Danielle Keats Citron, *Technological Due Process*, Washington University Law Review (2008), https://openscholarship.wustl.edu/cgi/viewcontent.cgi?article=1166&context=law_lawreview; see also Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, Emory Law Journal (April 3, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3553590.

¹⁸ For instance, evidentiary standards for admitting AI evidence in court have yet to be developed and are not encompassed in current *Daubert* standards guidance.

¹⁹ DHS's Artificial Intelligence Strategy, dated December 2020, includes the establishment of a DHS enterprise-wide AI Coordination and Advisory Council composed of internal subject matter experts to monitor and support the adoption of AI technology by DHS Components. See *U.S. Department of Homeland Security Artificial Intelligence Strategy*, U.S. Department of Homeland Security at 10 (Dec. 3, 2020), <https://www.dhs.gov/publication/us-department-homeland-security-artificial-intelligence-strategy>.

Blueprint for Action: Chapter 8 - Endnotes

²⁰ For issues relevant to AI system audits, see *Global Perspectives and Insights: The IIA's Artificial Intelligence Auditing Framework Part*, Institute of Internal Auditors (2018), <https://na.theiia.org/periodicals/Public%20Documents/GPI-Artificial-Intelligence-Part-II.pdf>.

²¹ See e.g., Audit Map (last accessed Jan. 3, 2021), <https://auditmap.ai/>; *The Next Generation of Internal Auditing—Are You Ready?*, Protiviti (2018), <https://www.protiviti.com/sites/default/files/united-states/insights/next-generation-internal-audit.pdf>.

²² See e.g., Bernhard Babel, et al., *Derisking Machine Learning and Artificial Intelligence*, McKinsey & Company (Feb. 19, 2019), <https://www.mckinsey.com/business-functions/risk/our-insights/derisking-machine-learning-and-artificial-intelligence>; Saqib Aziz & Michael Dowling, *Machine Learning and AI for Risk Management*, *Disrupting Finance* at 33-50 (Dec. 7, 2018), https://link.springer.com/chapter/10.1007/978-3-030-02330-0_3.

²³ Xuning (Mike) Tang & Yihua Astle, *The Impact of Deep Learning on Anomaly Detection*, Law.com (Aug. 10, 2020), <https://www.law.com/legaltechnews/2020/08/10/the-impact-of-deep-learning-on-anomaly-detection/>.

²⁴ Examples include baseline AI standards and policy guidance for biometric identification technologies; for government procurement of commercial AI products; and for federal data privacy standards.

²⁵ In the FY2021 NDAA, Congress directed the Secretary of Commerce, in consultation with other senior Executive branch officials, to establish the National AI Advisory Committee, including a Subcommittee on AI and Law Enforcement. The Subcommittee is tasked to “provide advice to the President on matters relating to the development of artificial intelligence relating to law enforcement.” Pub. L. 116-283, sec. 5104 William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

²⁶ These should seek to encourage contracts with companies that have transparent policies and practices in support of traceability and auditability and those that share information about how their technology works and how it performs in independent testing.

²⁷ “Federal government acquisition regulations require that agencies procure software commercially off-the-shelf whenever possible, due to their cost effectiveness. Only when no comparable systems exist are agencies permitted to develop government off-the-shelf solutions.” See Frances Duffy, *Ethical Considerations for Use of Commercial AI*, Johns Hopkins Applied Physics Laboratory at S-1 (December 2020). As standards and requirements for system development and testing evolve, it may be helpful for the government to “establish and maintain a list of COTS AI technologies that have been vetted and approved for micro-purchasing, based on their consistency with government security and testing standards, as well as their transparency.” This could facilitate both rapid procurement and proper assessment of a vendor’s consistency with Responsible AI practices. See Frances Duffy, *Supplement to Ethical Considerations for Commercial Use of AI: Implications of Acquisition Scale*, Johns Hopkins Applied Physics Laboratory (forthcoming).

²⁸ For example, policymakers and legislators will need to direct future attention to policies to preserve PCL as technological capabilities for ubiquitous sensing grow, e.g., in smart cities. In the future, ubiquitous sensing may make it impossible to distinguish U.S. persons’ data versus non-U.S. persons’ data for AI analytics. Another example for continued consideration includes the role of AI in filtering to remove U.S. persons’ information from bulk data and conversely using AI to reveal such information, as minimization and de-minimization guidance may evolve based on AI efficacy relative to the status quo.

²⁹ Disallowed outcomes and guidance will need to be updated over time as community norms and technical capabilities change.

³⁰ See, for example, *Remarks of Commissioner Rebecca Kelly Slaughter: Algorithms and Economic Justice*, FTC (Jan. 24, 2020) https://www.ftc.gov/system/files/documents/public_statements/1564883/remarks_of_commissioner_rebecca_kelly_slaughter_on_algorithmic_and_economic_justice_01-24-2020.pdf; *Artificial Intelligence and Machine Learning in Software as a Medical Device*, U.S. Food and Drug Administration (January 2021), <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device>.

³¹ See Byron Tau, *Homeland Security Watchdog to Probe Department's Use of Phone Location Data*, Wall Street Journal (Dec. 2, 2020), <https://www.wsj.com/articles/homeland-security-watchdog-to-probe-departments-use-of-phone-location-data-11606910402> (reporting that “DHS’s general counsel began examining [the agency’s use of location tracking data] after concerns were raised by several offices within the department that use of the technology wasn’t compatible with [Carpenter],” and that the DHS IG planned to investigate the matter).

³² In ODNI Director Avril D. Haines’ confirmation hearing, she was asked about the IC’s use of commercially available location data. She testified that she would “try to publicize, essentially, a framework that helps people understand the circumstances under which we do that and the legal basis that we do that under. . . . I think that’s part of what’s critical to promoting transparency generally so that people have an understanding of the guidelines under which the intelligence community operates.” Charlie Savage, *Intelligence Analysts Use U.S. Smartphone Location Data Without Warrants, Memo Says*, New York Times (Jan. 22, 2021), <https://www.nytimes.com/2021/01/22/us/politics/dia-surveillance-data.html>.

³³ Investigative reporting and opinion pieces have underscored the national security threats involved with smartphone location data. Charlie Warzel & Stuart A. Thompson, *They Stormed the Capitol. Their Apps Tracked Them*, New York Times (Feb. 5, 2021), <https://www.nytimes.com/2021/02/05/opinion/capitol-attack-cellphone-data.html?referringSource=articleShare>; Stuart A. Thompson & Charlie Warzel, *How to Track President Trump*, (Dec. 20, 2019), <https://www.nytimes.com/interactive/2019/12/20/opinion/location-data-national-security.html>; Stuart A. Thompson & Charlie Warzel, *Twelve Million Phones, One Dataset, Zero Privacy*, New York Times (Dec. 19, 2019), <https://www.nytimes.com/interactive/2019/12/19/opinion/location-tracking-cell-phone.html>.

³⁴ See Larry Magid, *IBM, Microsoft And Amazon Not Letting Police Use Their Facial Recognition Technology* (June 12, 2020), <https://www.forbes.com/sites/larrymagid/2020/06/12/ibm-microsoft-and-amazon-not-letting-police-use-their-facial-recognition-technology/?sh=34b473dc1887>; Asa Fitch, *Microsoft Pledges Not to Sell Facial-Recognition Tools to Police Absent National Rules*, Wall Street Journal (June 11, 2020), <https://www.wsj.com/articles/microsoft-pledges-not-to-sell-facial-recognition-technology-to-police-absent-national-rules-11591895282>.

³⁵ See *Ban Facial Recognition*, Fight for the Future (last accessed Feb. 4, 2021), <https://www.banfacialrecognition.com/map/>.

³⁶ The Department of Defense, the Drug Enforcement Administration, Immigrations and Customs Enforcement, the Internal Revenue Service, the Social Security Administration, the U.S. Air Force Office of Special Investigations, and the U.S. Marshals Service have all had access to one or more state or local face recognition systems. See Clare Garvie, et al., *The Perpetual Line-up: Unregulated Police Face Recognition in America*, Georgetown Law Center on Privacy & Technology (Oct. 18, 2016), <https://www.perpetuallineup.org/>.

³⁷ Such types of identification aided by AI include voice recognition and gait detection. An example of additional risks includes when biometric identification is coupled with other advancing capabilities; for instance, for identity recognition or for emotion recognition. See *Emotional Entanglement: China's Emotion Recognition Market and its Implications for Human Rights*, Article 19 (January 2021), <https://www.article19.org/wp-content/uploads/2021/01/ER-Tech-China-Report.pdf>. See also Drew Harwell & Eva Dou, *Huawei Tested AI Software that Could Recognize Uighur Minorities and Alert Police*, Report Says, Washington Post (Dec. 8, 2020), <https://www.washingtonpost.com/technology/2020/12/08/huawei-tested-ai-software-that-could-recognize-uighur-minorities-alert-police-report-says/>; Parmy Olson, *The Quiet Growth of Race Detection Software Sparks Concerns Over Bias*, Wall Street Journal (Aug. 14, 2020), <https://www.wsj.com/articles/the-quiet-growth-of-race-detection-software-sparks-concerns-over-bias-11597378154>.

Blueprints for Action

PART TWO



Chapter 9: A Strategy for Competition and Cooperation	413
Chapter 9 Annex: A Strategy for Competition and Cooperation	415
Chapter 10: The Talent Competition	421
Chapter 11: Accelerating AI Innovation	435
Chapter 12: Intellectual Property	465
Chapter 13: Microelectronics	483
Chapter 13 Annex: Executive Order on Microelectronic Strategy	489
Chapter 14: Technology Protection	493
Chapter 14 Annex: Technology Protection	511
Chapter 15: A Favorable International Technology Order	517
Chapter 15 Annex: A Favorable International Technology Order	559
Chapter 16: Associated Technologies	581

The following Blueprints for Action cover Part II of NSCAI’s Final Report. Part II, “Winning the Technology Competition” (Chapters 9-16), outlines AI’s role in a broader technology competition and recommends actions the government must take to promote AI innovation to improve all facets of national competitiveness and protect critical U.S. advantages. These Blueprints for Action complement the Commission’s Final Report and mirror its organizational structure.

Building upon the top-line recommendations in the Commission’s Final Report, the Blueprints for Action serve as more detailed roadmaps for Executive and Legislative branch actions to retain America’s AI leadership position. The Blueprints for Action identify who should take a particular action—Congress, the White House, or an Executive Branch department or agency. The Commission provides estimated increases in funding or appropriations as part of its recommendations. All recommendations that include funding figures should be considered estimates for consideration by Congress and/or the Executive Branch.

Chapter 9: A Strategy for Competition and Cooperation

Blueprint for Action

The United States should advance a comprehensive policy on China that promotes and protects a rules-based international order. By investing in U.S. competitiveness and resilience at home, safeguarding critical technologies, and deepening coordination with allies and partners, the United States can pursue cooperation with China—where it is in the national interest and from a position of strength. Properly sequenced and resourced, such a strategy would generate solutions to global challenges and leverage formal diplomatic dialogue to address critical issues around emerging technology.

Recommendation

Recommendation: Establish a High-Level U.S.-China Comprehensive Science and Technology Dialogue (CSTD)

The United States should establish a regular, high-level technology dialogue with China that benefits the American people, remains faithful to our allies, and presses China to abide by international rules and norms. The dialogue should focus on challenges presented by emerging technologies—to include AI, biotechnology, and other technologies as agreed by both sides. The CSTD should have two overarching objectives:

- Identify targeted areas of cooperation on emerging technologies to solve global challenges such as climate change, public health, and natural disasters; and
- Provide a forum to air a discrete set of concerns or friction points around specific uses of emerging technologies while building relationships and establishing process between the two nations.

The United States should be clear-eyed that the dialogue will not solve all our differences with China. The CSTD should be results-oriented, and it should achieve concrete outcomes for the American people.

Actions for the White House and the Department of State:

- **Establish the CSTD.**
 - Emerging technologies play an instrumental role in the economic, social, and security dynamics between the United States and China. Therefore, the CSTD should be established as part of a comprehensive strategy toward China that mobilizes democratic allies and partners in support of a rules-based international order.

- o The Department of State—in close coordination with the Office of Science and Technology Policy—should lead the CSTD.
- o The Department of State should build a process that is results oriented and aims to address challenges and opportunities in the current relationship between the United States and China related to the emerging technologies. For example:
 1. The CSTD should explore collaborative technological solutions to global challenges (e.g., climate change, healthcare and biodata, food safety and security, and natural disasters).
 2. The CSTD should identify areas of current challenges related to emerging technologies (e.g., data sharing and privacy, supply chain risk management, international standards and norms, and intellectual property) and develop a clear roadmap with milestones to address these issues.
- o The CSTD should initiate personnel exchanges and data-sharing frameworks to support and foster identified research projects with reciprocal access to information that can lead to concrete results.
- o The United States should identify leads for each of these topics (e.g., the Department of Energy, the U.S. Special Presidential Envoy for Climate, and the National Oceanic and Atmospheric Administration for climate change; the National Institutes of Health for healthcare; the U.S. Food and Drug Administration for food safety; and the Department of Defense and U.S. Agency for International Development for natural disasters).
- **Relation to strategic dialogue.** On a separate track from this CSTD, the Commission has recommended that the United States and Chinese governments convene a Strategic Security Dialogue (SSD) focused on eliminating misunderstandings and misperceptions on key strategic issues and threats and reducing the likelihood of inadvertent escalation. China has resisted U.S. attempts to create such a dialogue for nearly a decade, but its creation has never been more critical. The Commission's vision regarding the role of the SSD is explored in greater detail in Chapter 4 of this report.
 - o This dialogue should be the primary forum for discussions regarding practices surrounding AI-enabled and autonomous weapon systems and should include discussions on testing, doctrine, and use, and potentially the exploration of practical concrete confidence-building measures to mitigate risks.
 - o It is important to separate the SSD from the CSTD to ensure discussions related to conflict escalation and crisis stability are insulated from political forces which influence the broader U.S.-China bilateral relationship.

Chapter 9 Annex: A Strategy for Competition and Cooperation

Draft Executive Order Establishing the Technology Competitiveness Council

By the authority vested in me as President by the Constitution and laws of the United States of America, and in order to provide a coordinated process for developing technology policy and a national technology strategy and for monitoring its implementation, it is hereby ordered as follows:

Section 1. Policy. The national security, economic competitiveness, and domestic prosperity of the United States require a comprehensive and coordinated approach by the Federal Government to ensure long-term U.S. leadership across the entire suite of critical and emerging technologies. To achieve this objective, this order establishes a Technology Competitiveness Council to develop a National Technology Strategy and to coordinate policies regarding critical and emerging technologies across the Federal Government.

Section. 2. The Technology Competitiveness Council.

(a) *Establishment.* There is established a Technology Competitiveness Council (Council).

(b) *Membership.* The Council shall be composed of the following members:

- (i) the Vice President, who shall be Chair of the Council;
- (ii) the Secretary of State;
- (iii) the Secretary of the Treasury;
- (iv) the Secretary of Defense;
- (v) the Attorney General;
- (vi) the Secretary of Commerce;
- (vii) the Secretary of Energy;
- (viii) the Secretary of Homeland Security;
- (ix) the Director of the Office of Management and Budget;
- (x) the Assistant to the President for Technology Competitiveness;
- (xi) the Assistant to the President for National Security Affairs;

(xii) the Assistant to the President for Science and Technology;

(xiii) the Assistant to the President for Economic Policy;

(xiv) the Assistant to the President for Domestic Policy;

(xv) the United States Trade Representative;

(xvi) the Chairman of the Joint Chiefs of Staff; and

(xvii) the heads of other executive departments and agencies and other senior officials within the Executive Office of the President, as determined by the Chair.

A member of the Council may designate, to perform the Council functions of the member, a senior-level official who is part of the member's department, agency, or office and who is a full-time officer or employee of the Federal Government.

(c) Responsibilities of the Chair.

(i) The Chair or, upon his or her direction, the Assistant to the President for Technology Competitiveness, shall convene and preside over meetings of the Council and shall determine the agenda for the Council.

(ii) The Chair shall authorize the establishment of such committees of the Council, including an executive committee, and of such working groups, composed of senior designees of the Council members and of other officials invited to participate in Council meetings, as he or she deems necessary or appropriate for the efficient conduct of Council functions.

(iii) The Chair shall report to the President on the activities and recommendations of the Council. The Chair shall advise the Council as appropriate regarding the President's directions with respect to the Council's activities and national technology policy generally.

(d) Administration.

(i) The Council shall have a staff, headed by the Assistant to the President for Technology Competitiveness.

(ii) The Office of Administration in the Executive Office of the President shall provide the Council with such personnel, funding, and administrative support, to the extent permitted by law and subject to the availability of appropriations, as

directed by the Chair or, upon the Chair's direction, the Assistant to the President for Technology Competitiveness, to carry out the provisions of this order.

(iii) To the extent practicable and permitted by law, including the Economy Act, and within existing appropriations, agencies serving on the Council shall make resources, including but not limited to personnel and office support, available to the Council as reasonably requested by the Chair or, upon the Chair's direction, the Assistant to the President for Technology Competitiveness.

(iv) The heads of agencies shall provide, as appropriate and to the extent permitted by law, such assistance and information to the Council as the Chair may request to implement this order.

(v) Members of the Council shall ensure that their departments and agencies cooperate with the Council and provide such assistance, information, and advice to the Council as the Council may request, to the extent permitted by law.

(vi) The creation and operation of the Council shall not interfere with existing lines of authority and responsibilities in the departments and agencies.

(vii) On technology policy and strategy matters relating primarily to national security, the Council shall coordinate with the National Security Council (NSC) through the Deputy National Security Advisor for Cyber and Emerging Technology to create policies and procedures for the Council that respect the responsibilities and authorities of the NSC under existing law.

Section. 3. Functions of the Council. The Council shall:

(a) develop recommendations for the President on U.S. technology competitiveness and technology-related issues, advise and assist the President in development and implementation of national technology policy and strategy, and perform such other duties as the President may prescribe;

(b) develop and oversee the implementation of a National Technology Strategy as required by section 4 of this order;

(c) serve as a forum for balancing national security, economic, and technology considerations of U.S. departments and agencies as they pertain to technology research, development, commercial interests, and national security applications;

(d) coordinate policies across U.S. departments and agencies related to U.S.

competitiveness in critical and emerging technologies and ensure that policies designed to promote U.S. leadership and protect existing competitive advantages are integrated and mutually reinforcing; and

(e) synchronize budgets and strategies, in consultation with the Director of the Office of Management and Budget, in accordance with the National Technology Strategy.

Section. 4. National Technology Strategy. It is the policy of the United States to retain leadership in critical and emerging technologies essential to U.S. national security and economic prosperity. Within one year of the date of this order, and annually thereafter, the Council shall submit to the President a National Technology Strategy containing the following elements:

(a) an assessment of the U.S. Government's efforts to preserve U.S. leadership in key emerging technologies and prevent U.S. strategic competitors from leveraging advanced technologies to gain strategic military or economic advantages over the United States;

(b) a review of existing U.S. Government technology policy, including long-range goals;

(c) an analysis of technology trends and assessment of the relative competitiveness of U.S. technology sectors in relation to strategic competitors;

(d) identification of sectors critical for the long-term resilience of U.S. innovation leadership across design, manufacturing, supply chains, and markets;

(e) recommendations for domestic policy incentives to sustain an innovation economy and develop specific, high-cost sectors necessary for long-term national security ends;

(f) recommendations for policies to protect U.S. and allied leadership in critical areas through targeted export controls, investment screening, and counterintelligence activities;

(g) identification of priority domestic R&D areas critical to national security and necessary to sustain U.S. leadership, and directing funding to fill gaps in basic and applied research where the private sector does not focus;

(h) recommendations for talent programs to grow U.S. talent in key critical and emerging technologies and enhance the ability of the Federal Government to recruit and retain individuals with critical skills into Federal service; and

(i) methods to foster the development of international partnerships to reinforce domestic policy actions, build new markets, engage in collaborative research, and create an international environment that reflects U.S. values and protects U.S. interests.

Section. 5. Advisory Committee on Technology Competitiveness.

(a) There is established an Advisory Committee on Technology Competitiveness (Committee) to provide advice and recommendations to the Council and matters within the scope of the Council's responsibilities.

(b) The Committee shall include the Assistant to the President for Technology Competitiveness and not more than 16 additional members appointed by the President. The additional members shall include distinguished individuals from sectors outside of the Federal Government. They shall have diverse backgrounds and expertise in national security, economic competitiveness, and critical and emerging technologies relevant to the National Technology Strategy. The Assistant to the President for Technology Competitiveness, along with one non-Federal member of the Committee, shall serve as Co-Chairs. Members of the Committee shall serve without any compensation for their work on the Committee, but they may receive travel expenses, including per diem in lieu of subsistence, as authorized by law for persons serving intermittently in the government service (5 U.S.C. 5701-5707).

(c) The Committee shall meet as directed by the Co-Chairs of the Council and shall provide advice or work product solely for use by the Council in the performance of its duties under this order.

(d) The Office of Administration in the Executive Office of the President shall provide such funding and administrative and technical support as the Committee may require.

(e) The Committee shall terminate two years from the date of this order unless extended by the President.

Section. 6. General Provisions.

(a) If any provision of this order or the application of such provision is held to be invalid, the remainder of this order and other dissimilar applications of such provision shall not be affected.

(b) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.

(c) Nothing in this order shall be construed to impair or otherwise affect:

(i) the authority granted by law to an executive department or agency, or the head thereof; or

(ii) the functions of the Director of the Office of Management and Budget relating to budgetary, administrative, or legislative proposals.

(d) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.

Chapter 10: The Talent Competition

Blueprint for Action

The United States must dramatically invest in its artificial intelligence (AI) talent pipelines in order to remain at the forefront of AI now and into the future. It is imperative that the United States strategically invest in science, technology, engineering, and mathematics (STEM) education at all levels and improve the immigration system to allow for more AI talent to enter and remain in the United States. Therefore, this Blueprint for Action is organized into two broad categories of recommendations for strengthening the U.S. talent pipeline: the U.S. education system and immigration.

Talent Pipeline: U.S. Education System

Investments in STEM education are a necessary part of increasing American national power and improving national security. This requires the United States to reform its education system to produce both a higher quality and quantity of graduates.

Recommendation

Recommendation: Pass a New National Defense Education Act

In response to the Soviet launch of Sputnik in 1957, the United States passed the National Defense Education Act (NDEA) in 1958 to extend U.S. leadership in education and innovation.¹ The NDEA promoted the importance of science, mathematics, and foreign languages for students, authorizing more than \$1 billion toward decreasing student loans, funding for education at all levels, and funding for graduate fellowships. Many students were able to attend college because of this bill; 3.6 million students attended college in 1960, and by 1970, it was 7.5 million.² This act helped America win the Space Race and accelerated our ability to innovate, and it is widely regarded as one of the most successful pieces of education legislation in U.S. history.

Now is the time for a new NDEA. The NDEA greatly increased the number of Americans with a college degree, expanded the number of math and science teachers to meet the demand of the K-12 system after the postwar baby boom, and was focused on defense-centric fields, particularly a deficiency in mathematicians. The impacts of federal spending on higher education today are echoes of the investments made in the late 1950s by the Eisenhower administration. The United States needs a second NDEA (NDEA II) in order to address the current digital talent gap and prevent the United States from falling behind in the race for AI and STEM talent.

Actions for Congress:

- **Increase Funding for STEM- and AI-Focused After-School Programs**

- o STEM and AI-focused after-school learning programs expose students to STEM- and AI-related programs beyond normal school hours. The length of the school day limits teachers' ability to cover a myriad of topics. American elementary school students are exposed to an average of 20 minutes of science and 60 minutes of math during the school day.³ Given the short amount of time that teachers are able to spend on STEM in their classrooms, some school districts have begun to offer after-school programs that expose students to STEM in a less structured environment. More time spent studying STEM topics helps students' test scores, and for those who are underrepresented in STEM fields, federal funding for after-school programs will increase students' accessibility to quality educational tools.⁴ Appropriations for after-school programs should favor applications that are jointly submitted by a local educational agency and a community-based organization or other public or private entity as a way to defray costs and encourage community engagement.

- **Increase Funding for STEM- and AI-Focused Summer Learning Programs**

- o STEM- and AI-focused summer learning programs will encourage students to engage in STEM and AI activities during the months when students are typically unengaged and experience learning loss. The 21st Century Community Learning Centers Act is an example of a program that funds "academic enrichment opportunities during non-school hours for children, particularly students who attend high-poverty and low-performing schools" and has exhibited proven, positive results.⁵ Much like the after-school initiative, priority should be given to those applications that are jointly submitted by a local educational agency and a community-based organization or other public or private entity.

- **Allocate Funds for K-12 STEM Teacher Recruitment, Retention, and Training**

- o Teachers are an integral part of the learning experience for STEM subjects. One inequity is the lack of teachers with the requisite proficiency in STEM. Evidence shows that STEM teacher training for current teachers is sporadic, ineffective, and not effective in addressing the specific needs of individual students.⁶ Moreover, recruiting high-quality K-12 teachers with STEM experience and proficiency is difficult. This is particularly concerning, as teachers are one of the most influential aspects of school, having two to three times the impact of other components, such as leadership and school services.⁷ As the world continues to integrate technology into education, teachers must be taught how to use this technology as well as how to teach students the critical foundations and basic functions that come with it.⁸ Support should be given to school districts to create and execute teacher training in AI concepts, techniques, and curriculum design, with preference given to professional development courses that count against continuing education requirements for teacher certification.

- **Direct and Fund the National Science Foundation to Create STEM Scholarships and Fellowships**

- o We recommend that the NSF create 25,000 STEM undergraduate scholarships, 5,000 STEM PhD fellowships, and 500 postdoctoral positions over five years to increase the number and quality of STEM and AI practitioners that will reach the job market in a few years.⁹ Growing the nationwide STEM talent pool in high-demand

areas requires a pipeline of students who have studied relevant STEM coursework during their undergraduate careers. Between 2000 and 2017, the share of STEM bachelor's degrees earned—as a percentage of total bachelor's degrees earned in the U.S.—rose from 32% to 35%.¹⁰ The sharpest recent increases were among computer science and engineering majors.¹¹ For AI specifically, a degree in cognitive science or computer science with concentrations in AI or machine learning (ML) can pave the way for future careers in AI research or practice. AI is rarely offered as a major at the undergraduate level. Instead, universities offer standalone courses, a sequence of AI courses, or the option to study a technical major with a concentration in AI. Until a major in AI is more universally offered at U.S. universities, STEM scholarships will increase the number of individuals with the skills necessary to work on AI.

- o Scholarship and fellowship recipients should receive full tuition and room and board. Undergraduate recipients should receive a stipend of \$40,000 a year, and graduate recipients should receive a stipend of \$70,000 a year.¹² Combined with postdoctoral positions, this will bring the total cost to \$7.2 billion over five years.¹³

Actions for the Department of Education:

- **Add Elements of Computational Thinking and Statistics to Student Testing**

- o Computational thinking and statistics are vital for students to understand how AI works.¹⁴ As interdisciplinary fields, the use of computational thinking and statistics within AI can be found at all stages of discovery, from developing and planning studies to assessing the results. Critical thinking along with problem-solving are vital skills taught in statistics. Unfortunately, the majority of high schools in America do not require testing for skills related to computational thinking for graduation.¹⁵ There is no way to comprehensively measure U.S. students' overall abilities or aptitude for skills related to computational thinking and statistics. Students are taught what is needed to pass exams. Compared to other countries, many of which have statistics in their curriculum, the United States ranks low in math.¹⁶ By including subjects critical for computational thinking and statistics in standardized testing at the state level, the United States can gain a better understanding of students' capabilities and work to implement curriculum and lessons focused more on computational thinking and statistics in order to ensure students' success.

Recommendation

Recommendation: Require Statistics in Middle School and Computer Science Principles in High School

Actions for State Legislatures:

- **Require statistics as a required course in middle school and computer science principles in high school. Many fundamental concepts in AI, ML, and their subfields are applied statistics in disguise.¹⁷ The techniques and algorithms used are heavily based in statistical methods, such as cluster analysis and model selection. Statistics and computer science principles are needed to prepare students for AI courses, concentrations, and internships. Providing training in statistics starting in middle school will better prepare students for the increasingly advanced analytic techniques in demand for AI and STEM careers. Similarly, currently only 47% of U.S. high schools offer computer science coursework.¹⁸ This is much higher than just a decade ago, thanks to nationally organized initiatives, but this still leaves many high schools without computer science education.**

Moreover, adoption has been piecemeal and curriculum depth varies widely. Therefore, state action is needed.

- **On their own, neither statistics nor computer science are sufficient to teach students the concepts needed to understand AI. Having both allows students to experience the critical bases that must be covered early on in order to prepare students for a technological career. Simple math such as basic probability and summarizing numerical data is applying concepts of statistics and computer science.**

Talent Pipeline: Immigration

Immigration reform is imperative for strengthening the U.S. talent pipeline, particularly given the significant benefits the United States experiences due to highly skilled immigration. Therefore, the United States must pursue reforms to accelerate highly skilled immigration and retention of international students within the United States.

The following recommendations are intended to help the United States lead the world's development and implementation of AI by gaining a decisive majority of a critical and limited resource: AI talent. The recommendations will improve the United States' ability to attract talent to the United States and, just as important, away from competing countries.

The United States needs to take bold steps to ensure it wins the competition for international talent for years to come. Such steps should ensure that our immigration system attracts students, technical experts, and entrepreneurs; grants stability while they continue to contribute to the American economy and research environment; and retains students, entrepreneurs, and experts rather than sending them home or to competing countries. The best way to accomplish these goals and to send a clear message to AI and STEM talent around the world is to pass a National Security Immigration Act that specifically helps STEM talent remain in the United States, reduces the overall burden of the citizenship process, and creates specific paths for entrepreneurs.

Recommendation: Pass a National Security Immigration Act

1) Grant Green Cards to All Students Graduating with STEM PhDs from Accredited American Universities

This would issue an incredibly clear message to talented young people around the world that they are welcome in the United States and would ease their transition to American citizenship. It is a very aggressive maneuver to gain a larger share of the world's STEM talent.

Such a proposal is admittedly bold, but the benefits of attracting vetted, top-tier talent outweigh the risks. Bold measures are needed to preserve America's advantages in STEM fields today and to ensure we out-innovate and outperform competitors in the future.¹⁹ Few

Recommendation

other proposals are significant enough to make a dramatic difference in the competition for talent, or to force China into a dilemma on their domestic front. It is also noteworthy that similar proposals have received bipartisan support in the past.²⁰

Actions for Congress:

- **Amend 8 U.S.C. 1151(b)(1) to grant lawful permanent residence to any foreign national who:**
 - o Graduates from an accredited United States institution of higher education with a doctoral degree in a field related to science, technology, engineering, or mathematics in a residential or mixed residential and distance program;
 - o Has a job offer in a field related to science, technology, engineering, or mathematics; and
 - o Does not pose a national security risk to the United States.
- **Vetting for national security concerns should be enabled by the FBI and Intelligence Community**
- **Graduates granted lawful permanent residence through this program should not count against overall or country-of-origin green card caps**

2) Double the Number of Employment-Based Green Cards

Whether one aims for the United States to achieve AI dominance, grow gross domestic product (GDP), stimulate job growth, reduce government deficits, or bolster the solvency of the U.S. Social Security program, the most straightforward solution is the same: increase the number of highly skilled permanent residents. Under the current system, employment-based green cards are scarce: 140,000 per year, fewer than half of which go to the principal worker.²¹ This leaves many highly skilled workers unable to gain permanent residency and unable to transfer jobs or negotiate with employers as effectively as domestic workers. If underpaid, these workers cannot leave their jobs or bargain for better wages without risking revocation of the employer's green card sponsorship or even firing and forced departure from the United States. This decreases the appeal of joining the American workforce.

The H-1B system is problematic for most employers, as well, with a consistently oversubscribed "lottery" of 85,000 visas each year (of which 20,000 are reserved for advanced degree holders from U.S. universities).²² To reduce the backlog of highly skilled workers, the United States should double the number of employment-based green cards, with an emphasis on permanent residency for STEM and AI-related fields. If it were easier for U.S. employers to sponsor global talent for a green card as opposed to an H-1B visa, the H-1B program could then serve its originally intended function as a vehicle for truly temporary high-skilled work needs.

Action for Congress:

- **Amend 8 U.S.C. 1151(d)(1)(A) by changing “140,000, plus” to “280,000, plus”**

3) *Create an Entrepreneur Visa*

International doctoral students are more likely to want to found a company or become an employee at a startup than their native peers; but, in practice, they are less likely to pursue those paths. One reason is the constraints of the H-1B visa system.²³ Similarly, immigrant entrepreneurs without the capital to use the EB-5 route to permanent residency are forced to use other visas that are designed for academics and workers in existing companies, not entrepreneurs.²⁴ All of these issues make the United States less attractive for international talent and, just as important, reduce the ability of startups and other small companies, the main source of new jobs for Americans, to hire highly skilled immigrants that have been shown to improve the odds that the business will succeed.

Actions for Congress:

- **Create an entrepreneur visa. This visa should serve as an alternative to employee-sponsored, investor, or student visas and should instead target promising potential founders. Legislation should:**
 - o Define an entrepreneur as an alien whose organization and operation of a business would provide significant public benefit to the United States if allowed to stay in the country for a limited trial period to grow a company.
 - o Prioritize entrepreneurs active in high-priority fields such as AI or in fields that use AI for other applications, such as agriculture. The National Science Foundation should update the list of high-priority fields every three years.
 - o Use capital capture as a screening criterion for entrepreneurs.
 - o Emphasize job creation for Americans—potentially emphasizing underserved regions or areas with high unemployment—as a core factor in the assessment of significant public benefit.

4) *Create an Emerging and Disruptive Technology Visa*

A new nonimmigrant visa designed to attract top technology talent in critical fields would allow universities and businesses that work on AI and other emerging technologies access to a greater pool of talent necessary to create cutting-edge research. It would also respond more flexibly to labor market demands as new technologies emerge. The effect would be to “revitalize our country’s research ecosystem, empower our country’s innovation economy, and ensure that the United States remains a world superpower in the coming decades.”²⁵

Action for Congress:

- **Create an emerging and disruptive technology visa that:**
 - o Requires the National Science Foundation to identify critical emerging and disruptive technologies every three years;
 - o Allows students, researchers, entrepreneurs, and technologists in applicable fields to apply; and
 - o Does not include emerging and disruptive technology visa holders in any other visa category cap.

Recommendation

Recommendation: Broaden the Scope of “Extraordinary” Talent to Make the O-1 Visa More Accessible and Emphasize AI Talent

The O-1 temporary worker visa is for people with extraordinary ability or achievement.²⁶ O-1 visas are valid for three years and can be renewed annually an unlimited number of times. There is also no limit on the number of visas issued per year. Currently, about 15,000 to 18,000 new O-1 visas are issued annually.²⁷ For these reasons, the O-1 visa is generally a more flexible visa category than the H-1B visa, which is, with some exceptions, capped in duration and number.²⁸

While O-1 visas provide many advantages, they are a poor fit for many highly skilled workers due to the uncertainty of their criteria and the administrative burden of the application and adjudication process. Adjudicators determine an applicant’s eligibility through subjective assessments of whether applicants received nationally recognized prizes, have been published in major outlets, have done original work of major significance, and meet other similar criteria. For the sciences and technology, this aligns largely with academic criteria such as publications in major outlets and is not well suited for people who excel in industry.

Actions for the U.S. Citizenship and Immigration Service (USCIS):

- **Issue new guidance with clear and broad standards for regulatory criteria, such as what counts as a major outlet, nationally recognized prize, or original work.**
 - o For example, if a publication in a top-five academic journal within a scientist’s field counts as a major outlet, many PhD graduates would likely qualify.
- **Initiate a regulatory process to decrease the threshold for eligibility for an O-1 visa, for example by reducing the number of criteria an applicant has to fulfill.**
 - o The current standard is three out of eight criteria.²⁹
- **Broaden criteria to better accept non-academic AI and STEM accomplishments.**
 - o For instance, some top-tier engineers have not earned an undergraduate degree or published major papers, instead focusing on developing and monetizing cutting-edge technology in the private sector. New criteria should make O-1 visas more accessible to this demographic.

Recommendation: Implement and Advertise the International Entrepreneur Rule

The International Entrepreneur Rule (IER) allows USCIS to grant a period of authorized stay to international entrepreneurs who demonstrate that “their stay in the United States would provide a significant public benefit through their business venture.”³⁰ The IER would be relatively easy for the Executive Branch to implement and is more directly tied to job creation than most other immigration proposals, making it more helpful to most Americans.

Action for the President:

- **An immediate executive action could announce the administration’s intention to use the IER to boost immigrant entrepreneurship, job creation for Americans, and economic growth.**

Actions for the USCIS:

- **Announce that USCIS will give priority to entrepreneurs active in high-priority STEM fields such as AI, or in fields that use AI for other applications, such as agriculture.**
- **Use capital capture as a screening criterion for entrepreneurs.**
- **Emphasize job creation for Americans—potentially emphasizing underserved regions or areas with high unemployment—as a core factor in its assessment of significant public benefit.**

Recommendation: Expand and Clarify Job Portability for Highly Skilled Workers

The Department of Homeland Security (DHS) published a final rule in November 2016 that made a number of reforms to improve temporary work visa programs, including some measure of relief for workers tethered to the employer sponsoring their green card petition during a potentially decades-long waiting period.³¹ The rule allows workers on H-1B, O-1, and other temporary work visas to obtain open-market work permits for a one-year renewable period under compelling circumstances. Compelling circumstances include:

- Serious illness or disability faced by the worker or his/her dependents,
- Employer retaliation against the worker,
- Other substantial harm to the worker, and
- Significant disruption to the employer.³²

The criteria for compelling circumstances are too limited and ambiguous. Expanding visa holders’ ability to obtain a work permit would allow for greater rates of entrepreneurship, tighter skill-matching with new employers, and for visa holders to negotiate compensation on a level playing field with domestic workers.

Actions for the USCIS:

- **Clarify when highly skilled, nonimmigrant workers are permitted to change jobs or employers;**
- **Increase job flexibility when an employer either withdraws their petition for an H-1B or goes out of business, is acquired, or downsizes; and**
- **Increase flexibility for H-1B workers seeking other H-1B employment.**

Recommendation

Recommendation: Recapture Green Cards Lost to Bureaucratic Error

Congress mandates annual caps on the number of green cards that may be issued to certain family-based immigrants (226,000) and employment-based immigrants (140,000).³³ Because federal agencies do not want to exceed the annual green card caps, they generally issue fewer green cards than they are allowed to. Due to this trend, as of 2009, the Federal Government had not issued more than 326,000 green cards.³⁴ The number today is likely higher, but DHS has not published updated statistics.

Actions for the Departments of Homeland Security and State:

- **Publish an annual report on the number of green cards lost due to bureaucratic error.**
- **Review whether existing authorities can be used to:**
 - o Issue lost green cards the subsequent year without counting against green card caps.
 - o Prioritize highly skilled immigrants who have waited the longest, followed by highly skilled immigrants with long projected wait times.
- **If existing authorities are insufficient, engage with Congress to recapture green cards lost to bureaucratic error through special legislation.**

Blueprint for Action: Chapter 10 - Endnotes

¹ Pub. L. 85-864.

² *Sputnik Spurs Passage of the National Defense Education Act*, U.S. Senate (last accessed Jan. 29, 2021), https://www.senate.gov/artandhistory/history/minute/Sputnik_Spurs_Passage_of_National_Defense_Education_Act.htm#:~:text=The%20National%20Defense%20Education%20Act%20of%201958%20became%20one%20of.and%20private%20colleges%20and%20universities.

³ *Highlights From the 2018 NSSME+*, The National Survey of Science and Mathematics Education at 17 (Jan. 2019), <http://horizon-research.com/NSSME/wp-content/uploads/2019/01/Highlights-from-2018-NSSME.pdf>. Additionally, almost half of Americans believe that students don't spend enough time during school hours on STEM subjects. Cary Funk & Kim Parker, *Most Americans Evaluate STEM Education as Middling Compared with Other Developed Nations*, Pew Research Center (Jan. 9, 2018), <https://www.pewsocialtrends.org/2018/01/09/5-most-americans-evaluate-stem-education-as-middling-compared-with-other-developed-nations/>.

⁴ Kristen A. Malzahn, et al., *Are All Students Getting Equal Access to High-Quality Mathematics Education? Data From the 2018 NSSME+*, The National Survey of Science and Mathematics Education at 15 (Feb. 2020), <http://horizon-research.com/NSSME/wp-content/uploads/2020/02/Math-Equity-Report.pdf>.

⁵ *21st Century Learning Centers*, Department of Education (last accessed Jan. 1, 2021), <https://www2.ed.gov/programs/21stcclc/index.html>.

⁶ *Successful K-12 STEM Education*, National Research Council at 20-21 (2011), <https://www.nap.edu/catalog/13158/successful-k-12-stem-education-identifying-effective-approaches-in-science>.

⁷ Isaac M. Opper, *Teachers Matter: Understanding Teachers' Impact on Student Achievement*, RAND (2019), <https://www.rand.org/education-and-labor/projects/measuring-teacher-effectiveness/teachers-matter.html>.

⁸ Amy Johnson, et al., *Challenges and Solutions When Using Technologies in the Classroom*, Adaptive Educational Technologies for Literacy Instruction (2016), <https://files.eric.ed.gov/fulltext/ED577147.pdf>.

⁹ James Manyika & William H. McRaven, *Innovation and National Security: Keeping our Edge*, Council on Foreign Relations (Sept. 2019), <https://www.cfr.org/report/keeping-our-edge/recommendations/>.

¹⁰ Josh Trapani & Katherine Hale, *Trends in Undergraduate and Graduate S&E Degree Awards*, National Science Foundation at Figure 2-6 (Sept. 4, 2019), <https://nces.nsf.gov/pubs/nsb20197/trends-in-undergraduate-and-graduate-s-e-degree-awards>.

¹¹ *Id.*

¹² The \$70,000 stipend is intended to incentivize American students to pursue graduate research, rather than transitioning to the private sector directly after completing their undergraduate degree. Research has shown that higher stipends increase the number and quality of program applicants, likely "attract[ing] some potentially outstanding science and engineering students who would otherwise choose other careers." See Richard Freeman, et al., *Supporting "The Best and Brightest" in Science and Engineering: NSF Graduate Research Fellowships*, The National Bureau of Economic Research and Harvard University at abstract (Mar. 2006), https://users.nber.org/~sewp/Freeman_NSFstip_Proceedings.pdf.

¹³ Based on the Commission staff's research, the Commission calculates this total allotting an estimated \$175,000 per postdoctoral fellow per year.

¹⁴ Computational thinking can be defined as "a way of solving problems, designing systems, and understanding human behavior that draws on concepts fundamental to computer science." Center for Computational Thinking at Carnegie Mellon (last accessed Feb. 8, 2021), <http://www.cs.cmu.edu/~CompThink/>. Some current subjects relevant to computational thinking include computer science, coding, and statistics.

¹⁵ See *50 State Comparison: High-School Graduation Requirements*, Education Commission of the States (Feb. 2019), <https://internal-search.ecs.org/comparisons/high-school-graduation-requirements-01>. As shown in this 50-state comparison, unlike algebra, statistics is rarely listed as a graduation requirement. See *Id.*

- ¹⁶ Erin Richards, *Math Scores Stink in America. Other Countries Teach It Differently and See Higher Achievement*, USA Today (Feb. 29, 2020), <https://www.usatoday.com/story/news/education/2020/02/28/math-scores-high-school-lessons-freakonomics-pisa-algebra-geometry/4835742002/>.
- ¹⁷ Statistics includes foundations of probability, hypothesis testing, expected utility, decision analysis, and causality, and introductions to topics in the broader data sciences, such as basics of pattern recognition and machine learning.
- ¹⁸ *2020 State of Computer Science Education: Illuminating Disparities*, Code.org Advocacy Coalition, Computer Science Teachers Association & Expanding Computing Education Pathways Alliance (2020), https://advocacy.code.org/2020_state_of_cs.pdf.
- ¹⁹ According to the National Science Foundation (NSF), in 2018, 179,500 undergraduate and 233,600 graduate international students were enrolled in science and engineering programs in the United States. Beethika Kahn, et al., *The State of U.S. Science and Engineering 2020*, NSF (Jan. 15, 2020), <https://nces.nsf.gov/pubs/nsb20201/u-s-and-global-education#degree-awards>. It should not be assumed that all of these students would meet the listed criteria.
- ²⁰ A 2013 Senate-passed bill would have exempted all PhD and master's STEM degree holders (U.S. graduates) and all PhD holders in any field (worldwide graduates) from green card caps. Madeleine Sumption & Claire Bergeron, *Remaking the U.S. Green Card System: Legal Immigration Under the Border Security, Economic Opportunity, and Immigration Modernization Act of 2013*, Migration Policy Institute 2, at 8 (June 2013), <https://www.migrationpolicy.org/research/remaking-us-green-card-system-legal-immigration-economic-opportunity>.
- ²¹ William Kandel, *The Employment-Based Immigrant Backlog*, Congressional Research Service at 4-5 (March 26, 2020), <https://fas.org/sfp/crs/homesec/R46291.pdf>.
- ²² *H-1B Fiscal Year (FY) 2021 Cap Season*, U.S. Citizenship & Immigration Services (last accessed Jan. 4, 2021), <https://www.uscis.gov/working-in-the-united-states/temporary-workers/h-1b-specialty-occupations-and-fashion-models/h-1b-fiscal-year-fy-2021-cap-season>.
- ²³ Michael Roach, et al., *Are Foreign STEM PhDs More Entrepreneurial? Entrepreneurial Characteristics, Preferences and Employment Outcomes of Native and Foreign Science & Engineering PhD Students*, National Bureau of Economic Research at 12 (2019), <https://www.nber.org/papers/w26225>.
- ²⁴ William R. Kerr, *Global Talent and U.S. Immigration Policy: Working Paper 20-107*, Harvard Business School at 14 (2020), https://www.hbs.edu/faculty/Publication%20Files/20-107_0967f1ab-1d23-4d54-b5a1-c884234d9b31.pdf.
- ²⁵ Oren Etzioni, *What Trump's Executive Order on AI Is Missing: America Needs a Special Visa Program Aimed at Attracting More AI Experts and Specialists*, Wired (Feb. 13, 2019), <https://www.wired.com/story/what-trumps-executive-order-on-ai-is-missing/>.
- ²⁶ O-1A is the relevant O-1 category for STEM; it also encompasses those in "education, business, or athletics." *O-1 Visa: Individuals with Extraordinary Ability or Achievement*, U.S. Citizenship & Immigration Services (last accessed Jan. 29, 2021), <https://www.uscis.gov/working-in-the-united-states/temporary-workers/o-1-visa-individuals-with-extraordinary-ability-or-achievement>.
- ²⁷ *Nonimmigrant Visas Issued by Classification*, U.S. Department of State (last accessed Jan. 29, 2021), <https://travel.state.gov/content/dam/visas/Statistics/AnnualReports/FY2020AnnualReport/FY20AnnualReport-TableXVB.pdf>.
- ²⁸ *H1-B Fiscal Year (FY) 2021 Cap Season*, U.S. Citizenship & Immigration Services (last accessed Jan. 29, 2021), <https://www.uscis.gov/working-in-the-united-states/temporary-workers/h-1b-specialty-occupations-and-fashion-models/h-1b-fiscal-year-fy-2021-cap-season>.
- ²⁹ 8 C.F.R. 214.2(o)(3)(iii)(b).

Blueprint for Action: Chapter 10 - Endnotes

³⁰ *International Entrepreneur Parole*, U.S. Citizenship & Immigration Services (last accessed Jan. 29, 2021), <https://www.uscis.gov/humanitarian/humanitarian-parole/international-entrepreneur-parole>. There is currently no visa category well-suited to entrepreneurship in U.S. statutes related to immigration. The IER, which relies on parole authority, was initiated after legislative avenues were exhausted. Legislative fixes would be preferable but have so far proven politically infeasible.

³¹ 81 Fed. Reg. 82398, *Retention of EB-1, EB-2, and EB-3 Immigrant Workers and Program Improvements Affecting High-Skilled Nonimmigrant Workers*, U.S. Department of Homeland Security (Nov. 18, 2016), <https://www.federalregister.gov/d/2016-27540>.

³² *Id.*

³³ Julia Gelatt, *Explainer: How the U.S. Legal Immigration System Works*, Migration Policy Institute (April 2019), <https://www.migrationpolicy.org/content/explainer-how-us-legal-immigration-system-works>.

³⁴ A 2009 report to Congress indicates that some 242,000 unused family-based green cards were ultimately applied to the employment-based backlog. Congress also recaptured some 180,000 out of roughly 506,000 unused employment preference green cards via special legislation, leaving more than 326,000 green card numbers wasted out of the nearly 750,000 unused green cards. *Annual Report 2010*, Department of Homeland Security Citizenship and Immigration Services Ombudsman at 35 (June 30, 2010), https://www.dhs.gov/xlibrary/assets/cisomb_2010_annual_report_to_congress.pdf.

Chapter 11: Accelerating AI Innovation

Blueprint for Action

The United States remains the world's artificial intelligence (AI) leader. However, trends within the United States indicate underlying weaknesses. The Federal Government holds the responsibility to provide strategic direction and long-term resources to strengthen the nation's foundation for AI innovation. The United States—through government leadership, and in partnership with industry and academia—must increase the diversity, competitiveness, and accessibility of its AI innovation environment to ensure continued leadership.

Recommendation

Recommendation: Scale and Coordinate Federal AI R&D Funding

The United States must reinforce the foundation of technical leadership in AI by enacting a bold, sustained federal push to invest in AI R&D to foster a nationwide landscape of AI innovation and drive breakthroughs in the next generation of AI technologies by establishing a National Technology Foundation, funding AI R&D at compounding levels, establishing additional National AI Research Institutes, and making big bets on talent and innovative ideas.

Component 1: Establish a National Technology Foundation

In the wake of Russia's successful launch of the Sputnik satellite in 1957, Congress made significant investments in the National Science Foundation (NSF) to shore up U.S. leadership in science and technology.¹ Since then, the NSF has supported research across the frontiers of science and engineering, funding efforts that contributed to the development of the Internet, smartphones, and additive manufacturing.² However, in today's heightened geopolitical technology competition, even bolder action is needed to meet the promise of emerging and disruptive technologies like AI, drive U.S. innovation toward the national interest, and secure our economic future.

The Commission recommends the creation of a National Technology Foundation (NTF) as an independent federal agency and sister organization to the NSF to provide the means to move science more aggressively into engineering and scale innovative ideas into reality. This will require an organization that is structured to accept higher levels of risk and empowered to make big bets on innovative ideas and people. It also demands an emphasis on the transition of technology from the lab to the market.



□ NSCAI Recommendation

*Representing the current top 10 federal funders of non-defense AI R&D

The current federal R&D posture lacks an organization that provides the level of investment and focus in applied research and technology engineering commensurate with the benefit that technology breakthroughs could bring to the U.S. economy, society, and national security. In contrast to fundamental science, technology development embodies a more costly undertaking,³ requires the support of a diverse base of researchers and developers—including private-sector partners—and involves regular risk-taking. The Defense Advanced Research Projects Agency (DARPA) does this effectively, but for specific national security-focused ends and primarily through a prescribed program-based approach.

The NTF would drive technology progress at a national level by focusing on generating value at intermediate levels of technical maturity, prioritizing use-inspired concepts,⁴ establishing infrastructure for experimentation and testing, and supporting commercialization of successful outcomes. It would work in close concert with the NSF, DARPA, and other interagency partners to strengthen investment in domestic science and technology (S&T), providing the fuel for the development and delivery of AI and other technologies on which future economic progress and national security advantages rely.

To provide the level of attention to advance technologies of strategic importance, the NTF should focus efforts around a set of routinely updated priority research areas, such as those the Commission has identified as technologies critical to U.S. national competitiveness⁵:

- | | |
|--|-------------------------------|
| 1. Artificial Intelligence | 5. Robotics and Autonomy |
| 2. Biotechnology | 6. 5G and Advanced Networking |
| 3. Quantum Computing | 7. Advanced Manufacturing |
| 4. Semiconductors and
Advanced Hardware | 8. Energy Technology |

We do not underestimate the challenge of establishing a new institution; however, we see it as a strategic imperative. The NTF represents a long-term investment in America's ability to lead in AI and other disruptive technologies and apply technology toward efforts of societal importance. It would provide access to the resources and tools that could promote a national culture of experimentation and invention with new technology.

Given the criticality of holistically strengthening the national R&D landscape, the NTF should not detract from the level of appropriations for NSF, DARPA, or other existing federal R&D efforts. Rather, it should be instantiated as part of a broader approach that bolsters NSF as an institution of enduring, critical importance and amplifies federal support for technology R&D through existing channels as the organization gets off the ground.

Action for Congress:

- **Authorize and appropriate funding to support the establishment of the NTF.**
 - o To match the envisioned enlargement of U.S. technology efforts, federal investment

in the NTF should gradually increase from Fiscal Year 2022 to Fiscal Year 2026 for an ultimate estimated operating budget of \$20 billion per year.

- Additional funds for facilities and equipment necessary for the Foundation's creation, estimated at around \$30 million, should be made available starting in Fiscal Year 2022.
- o A National Technology Board—with members appointed by the President—should be created to provide policy direction to the NTF, supervise the Foundation's major initiatives, and ensure that its research focus areas are updated to reflect technology trends. The Board's directives and actions should be informed by the National Technology Strategy proposed by the Commission and, when necessary, coordinated with the Technology Competitiveness Council—both of which are separately recommended in this report.⁶
- o Jointly, a Director and Deputy Director appointed by the President should coordinate programming across the Foundation's directorates and with external organizations.
- o The NTF should be empowered to implement a portfolio of responsibilities:
 - Distribute funding through grants, cooperative agreements, and contracts awarded through competitive, risk-acceptant processes to academic and private-sector researchers, nonprofits, and consortia.
 - Manage a component of its funding through an innovation unit modeled on DARPA in which independent program managers would fund proposals from both industry and academia to advance solutions to forward-looking research questions.
 - Promote the transfer of technology advancements to the government as well as the commercial sector.
 - Run prize competitions to catalyze research around significant technology challenge problems.
 - Manage national technology resources and infrastructure that democratize an ability to build, test, and experiment.
 - Contribute to the success of the regional innovation clusters envisioned by the Commission by participating in the proposed technology program office and liaising with industry at Technology Research Centers.
 - Contribute to international R&D collaborations and standards-setting dialogues that strengthen U.S. strategic partnerships.

Component 2: Increase Federal Funding for Non-Defense AI R&D at Compounding Levels and Prioritize Key Areas of AI R&D

Research is the linchpin of America's global leadership in AI. However, current federal funding is not adequate to meet the growth of the field, let alone support its continued expansion.⁷ The Trump Administration's proposed budget for non-defense AI R&D in Fiscal Year 2021 was \$1.5 billion,⁸ a growth from around \$1 billion spent in Fiscal Year 2020.⁹ Further building on this investment, Congress included the National AI Initiative

Act of 2020 in the National Defense Authorization Act for Fiscal Year 2021, which creates a structure for a more strategic approach to harnessing AI and includes authorization for additional investments in AI at the NSF, Department of Energy (DoE), National Institute of Standards and Technology (NIST), and the National Oceanic and Atmospheric Administration (NOAA).¹⁰

National AI Initiative Act of 2020

- Created an executive branch entity within the Office of Science and Technology Policy to coordinate federal support for AI research and development, education and training, research infrastructure, and international engagement in order to achieve national priorities as defined in a regularly updated strategic plan for AI.¹¹
- Included provisions that established a National AI Research Resource task force, formalized the National AI Research Institute effort, and authorized funding for AI research at the National Science Foundation, the National Institute of Science and Technology, the Department of Energy, and the National Oceanic and Atmospheric Administration.

The government should build on these first moves and invest in AI R&D at compounding levels. Federal research funding holds the power to change the trends that are degrading the ability of the U.S. to continue to lead in AI, namely that academic research is weakening as a result of brain drain of professors and diversion of graduate students to industry, the domestic AI talent pipeline is not keeping up with government and industry needs, and national technical and ethical standards for development are lagging behind the technology.¹² Furthermore, federal support can spur the application of AI to other fields of science and engineering, which holds the potential for significant returns on investment.

Through sustained investments, federal support can serve to holistically strengthen AI R&D by embracing a range of initiatives—to include support for basic and applied research, shared research infrastructure, a network of AI R&D institutes, fellowships, and challenge competitions. Flowing investments through a diversity of agencies will create a vibrant fabric of funding, both mission-oriented and investigator-driven, that balances sustainment of evolutionary progress with big bets on revolutionary breakthroughs and supports innovation in academia and the private sector.

Actions for Congress:

- **Double annual non-defense AI R&D funding to reach \$32 billion by Fiscal Year 2026.**
 - o Congress should support compounding levels of federal funding for AI R&D, doubling investments annually from the baseline of \$1 billion in Fiscal Year 2020.
 - o Investments should be made across federal R&D funding agencies, notably the

proposed National Technology Foundation, DoE, NSF, the National Institutes of Health (NIH), NIST, and the National Aeronautics and Space Administration (NASA).

- o Significant funds should be appropriated to expand fellowship and scholarship programs.¹³ Augmented funding through these vehicles would support additional undergraduate and graduate students to pursue AI-related fields of study, helping to strengthen academia, grow the domestic talent pipeline, and provide pathways into government for technical talent. Similarly, career/faculty fellowship vehicles supporting researchers in academia would serve to stem the flow of researchers to industry and invest in top talent to pursue big ideas.

- **Commit to spending at least 1% of GDP on federally funded R&D.**

- o To maintain a strong base of innovation across S&T, Congress should pair AI-specific investments with an overall federal commitment to annually fund R&D at a level that reaches at least 1% of gross domestic product (GDP). This could be accomplished through steady growth over the next five years, at a rate of about \$15 billion per year.

Actions for the Office of Science and Technology Policy:

- **Balance Interagency AI R&D Investment Portfolios.**

- o The National AI Initiative should coordinate federal investments in AI R&D toward annual doubling benchmarks, through amplified research funding, fellowships, and establishment of research infrastructure.
- o The National AI Initiative should ensure that growth in funding occurs across multiple agencies and embodies a portfolio approach that leverages a diverse set of mechanisms, focused on a range of outcomes—advancement of basic science, solving specific challenge problems, and facilitating commercialization of breakthroughs.

- **Prioritize Critical AI Research Areas.**

- o Research investments should prioritize areas critical to advance AI technology that will underpin future national security and economic growth but may not receive significant private-sector investment, such as:
 - *Novel machine learning (ML) directions.* To further non-traditional approaches to supervised ML in an unsupervised or semi-supervised manner as well as the transfer of learning from one task or domain to another.¹⁴ Other directions include exploration of hybrid AI techniques that combine data-centric AI with different forms of model-based representations and inference methodologies to capitalize on complementary strengths.¹⁵
 - *Test and evaluation, verification and validation (TEVV) of AI systems.* To develop a better understanding of how to conduct TEVV and build checks and balances into the entire life cycle of an AI system,¹⁶ including improved methods to explore, predict, and control individual AI system behavior so that when AI systems are composed into systems-of-systems their interaction does not lead to unexpected negative outcomes. Understand context-specificity and degradation of performance in new and unseen environments.

- *Robust and resilient ML.* To cultivate more robust methods that can overcome adverse conditions and advance approaches that enable assessment of types and levels of vulnerability and immunity. Addressing challenges of multiple classes of adversarial ML attacks. Includes research on fairness.
- *Complex multi-agent scenarios.* To advance the understanding of interacting cohorts of AI systems, including research into adversarial vulnerabilities and mitigations, along with the application of game theory to varied and complex scenarios.
- *AI for modeling, simulation, and design.* To progress the use of rich simulations as a source of synthetic data and scenarios for training and testing AI systems, and to use AI to solve complex analytical problems and serve as a generative design engine in scientific discovery and engineering.
- *Advanced scene understanding.* To evolve perceptual models to incorporate multi-source and multi-modal information to support enhanced actionable awareness and insight across a range of complex, dynamic environments and scenarios.
- *Preservation of personal privacy.* To assure personal privacy of individuals is protected in the acquisition and use of data for AI system development and operation through advancements in anonymity techniques and privacy-preserving technologies such as homomorphic encryption, differential privacy techniques, and multi-party federated learning.
- *AI system risk assessment.* Advance capabilities to support risk assessment including standard methods and metrics for evaluating degrees of auditability, traceability, interpretability, explainability, and reliability.
- *Enhanced human-AI interaction and teaming.* To advance the understanding of human-AI teaming, including human-AI complementarity, methods for augmenting human reasoning abilities, and fluid handoffs in mixed-initiative systems. Also includes bolstering AI technologies to better perceive and understand human intention and communications, including comprehension of spoken speech, written text, and gestures. Advances in human-machine teaming will enable human interactions with AI-enabled systems to move from the current model of interaction where the human is the “operator” to a future in which humans have a “teammate” relationship with machines.
- *Autonomous AI systems.* To advance a system’s ability to accomplish goals independently, or with minimal supervision, from human operators in environments that are complex and unpredictable.
- *Toward more general AI.* Research persistent challenging problems and mysteries of human intellect, including ability to learn efficiently in an unsupervised manner; amass and apply commonsense knowledge; build causal models that provide robust explanations; exercise self-awareness, assessment, and control; and generalize and leverage knowledge learned about specific tasks to become proficient at another task.

Component 3: Triple the Number of National AI Research Institutes

NSF awarded grants for the first National AI Research Institutes in 2020, supporting seven university-based, multi-institution consortia organized around fundamental and applied areas of AI research—topics for which were determined through coordination with interagency and community stakeholders.¹⁷ NSF plans to fund a second round of institutes in 2021, coordinating support not only with interagency partners but also with private-sector stakeholders to launch eight additional institutes.¹⁸ Congress took steps to support the initiative through the National AI Initiative Act of 2020, which formalizes the effort, provides all agencies the authority to financially support formation of a National AI Research Institute, and directs NSF to bring together the institutes as an “Artificial Intelligence Leadership Network.”¹⁹

Expansion of this initiative would create a nationwide network of AI innovation that supports a breadth of AI research initiatives—advancing basic AI science, solving domain-specific challenges, and applying AI to other fields of science and engineering. Their establishment would increase training opportunities for students and research opportunities for academic faculty, national lab researchers, and non-profit research organizations; help grow the field outside of leading private universities and regional technology hubs; and strategically steer research toward areas that could advance the science of AI and applications that serve broader society and the national interest.

Action for Congress:

- **Direct and appropriate funds to expand the network of AI institutes.**
 - o Congress should direct and appropriate funds to NSF to expand the network of AI institutes three-fold over the course of the next three years—ideally resulting in a broad diversity of participating institutions, regions, and research concentrations.
 - o This investment would encompass 30 additional institutes, totaling \$600 million to sustain the additional institutes for the five-year duration of the grant awards. This would entail appropriations of \$200 million in Fiscal Year 2022, Fiscal Year 2023, and Fiscal Year 2024.

Action for the Office of Science and Technology Policy:

- **Integrate the network of institutes with national AI R&D infrastructure investments.**
 - o The National AI Initiative should ensure alignment of the National AI Research Institutes with strategic research priorities and integration with the national network of open AI test beds and the National AI Research Resource (see discussion of a National AI Research Infrastructure below).

Component 4: Invest in Talent that Will Transform the Field

Top talent in AI is a scarce commodity, and investing in talent holds the potential to not only unlock breakthroughs in the science and application of AI but also to attract and retain top talent in the United States.²⁰ Similarly, investing in research initiatives conducted by integrated, multidisciplinary teams is a proven mechanism to prompt breakthroughs, address complex problems, and challenge the status quo.²¹

The launch of an AI Innovator Award and complementary team-based AI award would strengthen the ability of federal AI research funding to push the boundaries of the field, providing a mechanism to complement ongoing investments in incremental progress with bets on revolutionary breakthroughs.

Actions for Congress:

- **Direct and fund establishment of an AI Innovator Award.**

- o Congress should direct and fund NSF to establish an AI Innovator Award, loosely modeled on the NIH Pioneer Award²² and the Howard Hughes Medical Institute Investigator Program²³ to create a mechanism that provides top researchers the flexibility to pursue big ideas without prescribed outcomes over the course of a five-year, renewable grant award.
 - Totalling around \$5.5 million per awardee for the five-year term, the awards would cover the full salary and benefits of the researchers at their respective institutions as well as a research budget that would support equipment and staff.²⁴
 - At its height, the program would support a maximum of 100 researchers at a time, reaching an annual funding level of around \$125 million for research support, with additional funds available for major equipment support.
 - Eligible researchers would be those at any career stage based at U.S. universities or research institutions who commit to spending 75% of their time on research.²⁵
 - Attention should be paid by the selection committee to the need for diversity among awardees in terms of gender, race, age, location, and primary focus area of study, as well as on the communication and leadership skills of applicants.
- o Congress should authorize NSF to:
 - Fund an external organization to administer the program.²⁶
 - Annually select 10 to 20 recipients for five-year, renewable terms and conduct selection through a small, rotating panel of AI experts.²⁷
 - Ensure selection of innovative candidates through an advocacy model process in which candidates are ranked in accordance with the maximum scores provided by reviewers, thereby placing priority on their upside potential.²⁸
 - Hold an annual meeting in which all awardees would share their work, providing a venue for meaningful feedback between review cycles and helping build a community of innovation among the top U.S.-based minds in AI.

- o Congress should require NSF to assess the program after seven years of operation to determine whether the program should continue to expand or operate at a lower number of awards and to evaluate the impact of the funding level and award term on the research conducted by participants.

- **Direct and fund establishment of a team-based AI research award.**

- o Congress should direct and fund NSF to work with the same external organization as the AI Innovator Award to create a team-based award to support bold, interdisciplinary research initiatives that apply AI to solve complex challenge problems or pursue use-inspired basic research efforts.
 - The program should begin with an annual budget of \$50 million, growing to a sustained annual budget of \$250 million by its fifth year of operation.
- o Congress should authorize the NSF to:
 - Fund an external organization to administer the program.
 - Select five to 10 teams annually for non-renewable, five-year terms, awarding \$4 million to \$10 million per year for the five-year term of the award.²⁹

Recommendation: Expand Access to AI Resources through a National AI Research Infrastructure

Recommendation

If not addressed, the growing divide between “haves” and “have nots” in AI R&D will degrade the long-term research and training functions performed by U.S. universities, limit the ability of small businesses to innovate, and exacerbate the lack of diversity in the field.³⁰ While developments in the past five years have dramatically increased access to baseline ML tools and cloud-based computation, progress on the cutting edge of many important AI approaches requires significant amounts of data and computing power, expensive infrastructure, and substantial hardware and software engineering.

The United States should foster the world’s leading environment for AI innovation through democratized access to AI R&D that supports more equitable growth of the field and expansion of AI expertise across the country; enables application of AI to a broad range of fields of science and engineering, commercial sectors, and public services; and fuels the next waves of innovation.

Component 1: Launch the National AI Research Resource

Since the explosion of deep learning in 2012 and accompanying growth in use of specialized hardware for AI computing, there has arisen what some have termed the “compute divide”—a disparity in access between large technology companies and elite universities and mid- and lower-tier universities to the resources necessary for cutting-edge AI research.³¹ Availability and type of compute resources have been found to levy “outsized” influence in the direction of research pursued by researchers, as has the ascendancy of the well-equipped firms in shifting the overall direction of AI research toward applied, “narrow AI” efforts.³²

To bridge the compute divide, the Federal Government should establish a National AI Research Resource (NAIRR) to provide verified researchers and students with access to compute resources, co-located with AI-ready government and non-government data sets, educational tools, and user support.³³ This infrastructure should leverage public-private partnerships and cutting-edge private-sector technology and build on existing government efforts³⁴—avoiding high startup costs of a government-run data center. Congress has taken the first step in the Fiscal Year 2021 National Defense Authorization Act, implementing a component of the Commission’s prior recommendation to create a task force to develop a roadmap for a NAIRR.³⁵ The result of this effort will be due to Congress 18 months after appointment of task force members.

Action for Congress:

- **Authorize and appropriate \$30 million for implementation of the NAIRR roadmap.**
 - o Congress should authorize and appropriate funds to immediately implement the roadmap developed by the NAIRR task force.
 - The resource should be sustained at an initial level of \$30 million annually, amplified by contributions from private-sector partners, and scaled as it matures and gains users.
 - Funding would support staffing of the program and the cloud resources, augmented through public-private partnerships. Staff would be responsible for maintaining and improving the architecture solution, curating data sets, building interfaces and tools, and providing support to researchers.

Component 2: Create a Network of National AI Testbeds to Serve the Academic and Industry Research Communities

Sponsored through various federal agencies, this network of national AI testbeds would provide real-world, domain-specific resources open to the academic, business, and government research communities to drive basic and applied research to address complex problems and develop robust, usable AI systems ripe for commercialization (for example, a self-driving vehicle test range, an instrumented humanitarian aid and disaster relief test site, or an instrumented home environment). Such resources would help establish and maintain benchmarking standards that enable measurable research progress through comparable approaches and reproducibility testing.

Testbeds should support experimentation with both novel software and hardware, equipped with rich simulation capabilities to model the physical world. Supported by simulated, live, and blended environments, these platforms would support research and experimentation that tackles open-ended, real-world problems. Furthermore, they should be architected to collect valuable data that could be made accessible to the community for training and evaluation, providing additional fuel for progress.

Action for the Office of Science and Technology Policy:

- **Coordinate agency investments in AI R&D testbed facilities.**
 - o The National AI Initiative should coordinate agency investments in AI testbed facilities through the annual budget process, aligning investments with research priorities issued in the initiative's strategic plan. Attention should focus on modernizing existing resources to support data-driven and AI-enabled technologies.³⁶

Action for Federal Agencies:

- **Invest in domain-specific AI R&D testbeds through upgraded or purpose-built facilities.**
 - o Investment in the suite of national AI testbeds should be made across multiple federal agencies, facilitating creation of domain-specific resources open to the broader research community. Focus areas of each testbed should be aligned with priority AI research areas and in support of existing federal AI investments.
 - o Testbeds should be set up as “user facilities” that maintain a hybrid approach of awarding grants for use and charging fees to those not selected for grant funding. User fees would assist in maintaining the testbeds and supplementing the amount of funding available for grants.

Action for Congress:

- **Support agency funding requests for establishment of AI R&D testbeds.**

Component 3: Invest in Large-Scale, Open Training Data

Data is critical currency for today's popular AI approaches. Promising work in the realm of low-shot learning, semi-supervised learning, and learning from synthetic data provides glimpses of a future in which performance of an AI system is not directly tied to big data, and the Federal Government should continue to prioritize funding for research in these areas. However, balancing these bets on the future with investments in resources to further U.S. leadership in the current leading AI approaches would strengthen the foundation of both current and future AI-based technology and applications.

Building AI systems and solutions for new domains and application areas relies on availability of specialized data that have been cleaned and organized for use. Federal support for well-designed, publicly-available data sets and provision of AI-ready government data sets would help drive research progress in AI and its application to other fields of study. Currently, a sizable amount of government data that is legal to share with trusted non-government researchers is not being shared due to a lack of confidence in cybersecurity and privacy-protecting technologies and a lack of willingness to accept risk.

Responsibly creating pipelines for the curation, hosting, and maintenance of complex data sets would set the foundation for future AI capabilities, help strategically steer the research community toward issues in the public interest, and advance technology around data set lifecycle maintenance.

These data investments could be further augmented by and created in support of the domain testbeds recommended above and hosted through the NAIRR. This integration could foster creation of data sets to support benchmarks within the testbeds as well as generate rich data from testing that could be provided back out to serve the research community. Access to resources should be granted to researchers with verified research efforts and governed by appropriate compliance controls based on the type of data and metadata contained in the data set.

Actions for the Executive Branch:

- **Issue a common policy and set of best practices.**
 - o Leveraging the work of NIST,³⁷ the U.S. Chief Data Officer should issue a common policy and set of best practices to support release of AI-ready government data to the public and work with industry and academia to adopt compatible policies and best practices for reciprocal sharing and documentation.
- **Provide incentives to industry and academia to make available select data sets.**
 - o The U.S. Chief Data Officer should develop incentives for industry and academia to make available select data sets on the NAIRR that would be managed and accessed alongside government-owned data sets.
- **Support NSF-funded cybersecurity and privacy researchers to make government data accessible for research purposes.**
 - o The National AI Initiative should coordinate NSF-funded cybersecurity and privacy researchers to undertake rotational assignments at federal agencies³⁸ and work closely with agency personnel and data stewards to responsibly unlock access to more of the government's data holdings for the purpose of stimulating AI research and innovation.
 - o Researchers would apply promising methodologies for protecting data and privacy in a controlled manner, providing a proving ground for new approaches and objective evidence to justify evolving data-sharing policies and practices. This could include creating secure environments for verified researchers to access more sensitive government-held data.

Actions for Congress:

- **Unlock public data for AI R&D.**
 - o Congress should fund teams of data engineers and data scientists organized through the U.S. Digital Service to unlock public data currently held by the government for use by the AI research community.³⁹

- o These teams would prioritize, clean, and curate non-sensitive public data sets to make them AI-ready and structure enduring processes to capture, clean, and regularly update data that would be hosted on a platform such as NAIRR, accessible by verified U.S.-based researchers.

- **Fund an AI data program at the Department of Energy.**

- o Congress should appropriate \$25 million⁴⁰ per year for the next five years to DoE to administer an AI data program that would create exemplar, complex data sets and maintain them as living, regularly updated resources. These could include specialized data sets in physical, biological, earth, and engineering sciences, as well as social sciences.⁴¹
- o The program should be coordinated through the National AI Initiative to ensure data sets created steer the research community in desired directions.
- o Congress should direct DoE to work closely with NIST to develop standards for the data—to include standards for documentation, data modeling, data engineering, and data formats—as well as to advance the methods and tools necessary to support the data lifecycle.

Component 4: Sponsor an Open Knowledge Network

Open knowledge networks (or repositories) with massive amounts of world knowledge could fuel the next wave of AI exploration, driving innovations from scientific research to the commercial sector. Today, only the biggest tech companies have the resources to develop significant knowledge graphs and networks.

Various federal agencies have invested in specialized, domain-specific knowledge networks that could provide a starting point for an open knowledge network.⁴² Beginning with a push to federate and map together existing specialized knowledge networks and government data platforms, and then building in real-world knowledge and context, the government could sponsor an Open Knowledge Network that would serve verified U.S.-based companies and researchers of all backgrounds to use world knowledge to develop AI systems that operate effectively and efficiently. This type of resource, particularly if paired with the complementary research infrastructure above, could unlock frontiers of technology yet unexplored.

Action for the Office of Science and Technology Policy:

- **Hold an innovation sprint to build an open knowledge network roadmap.**

- o Leveraging prior work undertaken through the Networking and Information Technology Research and Development (NITRD) program Big Data Interagency Working Group,⁴³ the Office of Science and Technology Policy should hold an innovation sprint to build a roadmap to establish an open knowledge network in a phased manner.

Action for Congress:

- **Direct and fund implementation and management of the open knowledge network.**
 - o Congress should direct and fund the NSF to implement and manage the open knowledge network, appropriating \$25 million per year for the next five years and encouraging NSF to leverage partnerships with industry stakeholders where possible.⁴⁴

Recommendation

Recommendation: Leverage Both Sides of the Public-Private Partnership

U.S. companies are at the forefront of AI R&D, and their investments benefit consumers globally through the rapid development and adoption of AI-enabled products. But the impact of AI-enabled products on U.S. society and national security has largely come as an afterthought. The speed of technology development by the private sector has vastly outpaced federal policies and regulations. To address these challenges, the public and private sector must share responsibility for the safety, security, and well-being of Americans. The following recommendations would make the government a better partner for industry, broaden the benefits of strategic emerging technologies like AI through regional innovation clusters, and expand opportunities to access AI research and education through private-sector philanthropy.

Component 1: Create Markets for AI and Other Strategic Technologies

The government's buying power cannot compete with a global consumer market, but it can influence investment decisions in technologies essential to overall U.S. technical leadership.⁴⁵ Many potential public-sector applications of AI, such as education and labor, fall under agencies with limited R&D budgets. As the government increases investment in basic research, it must also fully leverage its purchasing power to support AI and other strategic technologies.⁴⁶ The scale of government funding can influence the research priorities and viability of early-stage startups, which often succeed or fail in the first year; and, if leveraged collectively, it can draw private-sector resources toward areas of strategic priority. This makes investors and technology companies important partners for AI R&D that can build future defense and national security capabilities.

Yet the government remains a difficult customer—especially for small and medium-sized businesses—because of its complex contracting process and unique requirements. Making the U.S. government a more compelling customer and effective buyer of commercial technology will help drive technology development in the commercial sector that is in the national interest. It will also assist the government in almost every aspect of its mission, from providing basic public services to driving economic policy and protecting national security.

Actions for the General Services Administration (GSA):

- **Promote the application of AI across the U.S. Government.**
 - In fulfilling its mandate to facilitate the adoption of AI technologies in the Federal Government,⁴⁷ the AI Center of Excellence (AI CoE) should look first to readily available commercial off-the-shelf (COTS) technology that can be tailored for government use.
 - AI CoE should work with Federal technical leadership,⁴⁸ including the U.S. Chief Technology Officer, Chief Information Officer Council, and the National AI Initiative,⁴⁹ to identify government needs and opportunities and expedite the adoption of commercial AI applications across federal agencies.
 - The AI CoE should leverage existing digital governance efforts across the Executive Branch, including GSA's 18F and the U.S. Digital Service, and technical talent exchange programs, including GSA's Presidential Innovation Fellowship, to bring sufficient technical expertise and commercial proficiency to this effort.⁵⁰
- **Communicate federal AI capability priorities to the private sector.**
 - The AI CoE should add federal procurement priorities and agency capability needs to its publicly available website, which contains information regarding programs, pilots, and other initiatives.⁵¹

Actions for the U.S. Small Business Administration:

- **Publish a digital technology “playbook” for small businesses.**
 - A playbook for small businesses should outline paths for companies interested in doing business with the U.S. government and explain in a single place⁵² how to navigate challenges like obtaining access cards to government facilities. Such a resource would make the acquisitions process more transparent and reduce the need for companies to hire outside help.
 - The playbook should be developed and reviewed by personnel with technical and commercial proficiency, for example Presidential Innovation Fellows or staff from the U.S. Digital Service, and written in language that technology startups with no prior government experience can understand.
 - The playbook should be aggressively publicized to increase its visibility.
- **Bridge public and private investment through the Small Business Innovation Research (SBIR) Program.⁵³**
 - Support the efforts of participating federal agencies to modernize SBIR to more effectively develop and deploy AI solutions and encourage broader participation of American technology startup and small-business companies.
 - Expand pilot programs that offer supplemental funding to bridge the gap between current SBIR/Small Business Technology Transfer (STTR) Phase II awards and Phase III scaling efforts.⁵⁴

- Expand pilot programs that offer larger funding amounts⁵⁵ and private-sector matching opportunities to support higher technology readiness levels common in DoD SBIR contracts.⁵⁶
- o Update SBIR Policy Directive to allow programs to require matching private-sector funds as early as Phase II.⁵⁷

Actions for Department of Defense and Intelligence Community:

- **Adopt a “hoteling” model to allow small- and medium-sized technology companies to access classified facilities on a flexible basis.**
 - o The Digital Ecosystem described in Chapter 2 of this report would establish prototypical platform environments for contributors and users, including cleared personnel from AI companies. Flexible access to classified spaces would speed development cycles and help companies more regularly engage with current or potential customers within the national security enterprise, leading to more tailored and effective solutions delivered more quickly.
- **Simplify the contracting process to attract non-traditional vendors.**
 - o Amend the Defense Federal Acquisition Regulation to allow commercial performance to be considered more widely in the contracting process. The U.S. government can benefit from broader adoption of best-in-class commercial AI software.
 - o Allow for pilot use of commercially available digital application tools and access portals for SBIR and other non-traditional contracting vehicles.⁵⁸
- **Commit to growing the national security innovation base.**
 - o DoD should set a target of increasing its contracts with early-stage technology firms by four times over the five-year Future Years Defense Program.⁵⁹ This will also require growing the budgets of successful but nascent innovation initiatives such as the Defense Innovation Unit.⁶⁰
 - To this point, DoD has focused on a large number of small bets without following up with larger later-stage investments. Larger contracts for later-stage companies would help scale validated solutions that meet military requirements.
 - o The Under Secretary of Defense for Acquisition and Sustainment and the Service Acquisition Executives should encourage Acquisition Category programs of all sizes to solicit bids from at least one non-traditional contractor per program.
- **Strengthen return on SBIR investments.**
 - o Review, modernize, and streamline SBIR processes to encourage broader participation of American technology startup and small-business companies.⁶¹
 - Program officers should clearly communicate pathways to transition, including milestone criteria and dollar amounts, to SBIR awardees so that they can plan and resource accordingly.
 - Explicitly allow SBIR contracts to leverage any “color of money” as matching funds up to the amount of SBIR funding.

- o Enable successful prototypes to scale through sufficient funding, early access to customers and operators, and better due diligence on the commercialization prospects of a company.⁶²
 - Military Service and Office of the Secretary of Defense (OSD) SBIR programs should allocate a portion of SBIR funding for scaling successful SBIR projects through Phase II enhancements.⁶³
 - Program Offices should provide program dollars alongside matching SBIR funds to increase the likelihood of transition.
- o Continue efforts to align SBIR program with Department technology priorities to focus investments on subsets of key technologies on which private-sector R&D can help advance.⁶⁴
 - The Office of the Under Secretary of Defense for Research and Engineering should introduce a special solicitation on AI that invites solutions across a diversity of AI approaches⁶⁵ and a range of technology readiness levels.⁶⁶

Component 2: Form a Network of Regional Innovation Clusters Focused on Strategic Emerging Technologies

Competition is critical to a vibrant national security innovation base.⁶⁷ If a strategic industry lacks competition, one wrong bet by an incumbent can place the nation's technological leadership in jeopardy.⁶⁸ The U.S. government should create an environment in which innovative startups are able to disrupt inefficient or outdated ways of doing business and grow into industry leaders themselves. The right mix of policies and incentives can help firms overcome mounting barriers to entry at the cutting edge of emerging technologies like AI.⁶⁹ This approach will promote innovation in industries that are essential to U.S. leadership in AI and the nation's economic and technological competitiveness more broadly.⁷⁰

As the Commission noted in its *2019 Interim Report*, the clustering of technology firms in regions like Silicon Valley yields a more dynamic and globally competitive industry by expediting knowledge sharing and sharpening domestic rivalry.⁷¹ However, this trend has benefited some regions and demographics more than others.⁷² To spur regional innovation across a broader swath of the nation, the U.S. government should support the growth of technology clusters in regions with latent innovation potential. Broader in mission and scope than existing models within the U.S. government, such an initiative would democratize access to federal R&D resources so that small firms could compete in industries with high barriers to entry like AI. By facilitating the exchange of technology and talent between the public and private sectors, the U.S. government would also be well positioned to establish new contracts and intellectual property sharing agreements for commercial technologies that are critical to U.S. national security.

Actions for Congress:

- **Establish an interagency program office responsible for coordinating a network of regional innovation clusters focused on R&D and commercialization of strategic emerging technologies.**
 - o The program office should be hosted by the Department of Commerce at NIST and staffed by representatives from U.S. departments and agencies with experience in and missions related to strategic emerging technologies.⁷³ The program office should also draw on expertise from the private sector and academia through talent-exchange programs and external advisory arrangements.
 - o Congress should authorize \$5 million for the creation of the program office and task it with designating regional innovation clusters in qualified locations throughout the United States via a competitive process, as described below in detail. As a first step, the program office should solicit bids for financial assistance from applicants focused on the R&D and commercialization of strategic emerging technologies. In assessing bids, the program office should consider the following criteria:
 - *Location.* Clusters should be equitably distributed throughout the United States in regions with latent innovation potential, taking into account factors such as proximity to federal R&D facilities, the level of support from state and local governments, the presence of and value proposition for leading firms and research institutions, and the size and education level of the local workforce.⁷⁴
 - *Subject area.* Clusters should be organized around the research, development, and commercialization of strategic emerging technologies that are critical to U.S. national competitiveness. Of particular interest are technologies that enable advances in adjacent sectors and whose domestic production would directly benefit U.S. national security, such as microelectronics.⁷⁵
 - *Economic feasibility.* To maximize the impact of federal resources and ensure self-sustainability of the clusters, financial assistance should only be awarded to applicants that demonstrate the existence of a nascent cluster in their region.⁷⁶
 - o The program office should establish Technology Research Centers (TRCs) for each cluster to facilitate collaboration between participants. By forming sustained partnerships with anchor institutions, each TRC should strive to advance the research, development, and commercialization of strategic emerging technologies.⁷⁷
 - *Leverage talent.* TRCs should host researchers on temporary assignments from U.S. departments and agencies, establish talent exchanges with local firms and research institutions, and fund multi-year, postdoctoral fellowships for the commercialization of research.⁷⁸
 - *Encourage technology transfer.* TRCs should host program managers from U.S. departments and agencies responsible for transitioning basic research into commercially viable technologies, identifying national security use cases and end users within the U.S. government, and initiating new government contracts for those products.
 - *Generate intellectual property.* TRCs should establish intellectual property-sharing agreements with cluster participants to encourage government adoption of commercial technologies. When appropriate, research should be

published in the open-source domain to encourage advances in the broader science and technology community.

- *Bring government resources to bear.* TRCs should facilitate participants' access to federal computing resources, curated government data sets, testing infrastructure and ranges, and other R&D facilities at low cost.⁷⁹
 - o The program office should play a high-level coordination role that includes supervising the operation of TRCs, facilitating R&D collaboration between clusters, and promoting the commercialization of technologies with national security use cases.
- **Enact a package of provisions that incentivizes industry and academia to participate in clusters.**
 - o Provisions should include tax incentives to locate near the cluster, competitive research grants, loan guarantees, and seed funding. A complementary approach should be taken by state and local governments. These policies could be modeled on Opportunity Zones, which have stimulated investment in regional economies.⁸⁰
 - *Investment tax credits.* To compete with incentives offered by foreign countries, Congress should establish investment tax credits for firms participating in regional innovation clusters. While the details of these tax credits will vary by sector, one example is the investment tax credit for semiconductor manufacturing facilities and equipment proposed in Chapter 13 of this report.
 - **Provide funding to each cluster for at least five years, with matching investments from public- and private-sector partners.**
 - o Within one year, the program office should request from Congress the necessary funding for the designation of up to 10 clusters. This funding should be matched at least 1:1 by investment from private companies, state and local governments, and federal agencies, with a target of each cluster initially receiving a total of \$50 million annually. This annual amount should increase as demand and capacity at each cluster expands over time.⁸¹ These funds would be used to operate the TRCs, maintain R&D facilities, issue research grants, and seed startups.

Component 3: Establish a Private Sector-Led Competitiveness Consortium

The private sector shares responsibility with the government to strengthen the foundations of the R&D ecosystem that underpins breakthroughs they will commercialize and the training pipeline needed to meet their increasing demand for technical talent.

Companies are already struggling to find these qualified applicants for technical roles, with one estimate showing more than 400,000 open computing jobs nationwide.⁸² Furthermore, as described above, researchers in academia who will undertake the high-risk, high-gain research that will push the frontiers of the field are finding themselves locked out from the computing and data resources needed to fuel this work. How well the nation addresses this looming challenge has widespread implications for the economy, society, and U.S. global competitiveness.

Chapter 10 of this report describes in detail recommendations to revamp the U.S. educational system to equip Americans for the jobs of the future, and this chapter details the extensive investments the Federal Government should make in AI R&D. However, corporations should also consider their responsibility to prepare citizens for the future they are inventing and maintain the strong foundation of national innovation from which they benefit. Toward that end, many firms are already having a positive impact beyond their bottom lines through corporate social responsibility efforts. STEM education programs and job training feature prominently in the charitable-giving arms of leading tech firms.⁸³ Yet the scale of the challenge is too broad for individual firms to address in isolation, despite their generosity.

Actions for the Private Sector:

- **Donate \$1 billion over five years.**
 - o Providing every American an opportunity to increase their technical proficiency requires bold action from government, academia, and industry to coordinate, prioritize, and scale programs that broaden AI research opportunities and instill digital proficiency.⁸⁴ For the private sector to meet this call to action, the Commission calls upon industry to donate \$1 billion over the next five years to support AI education and upskilling and provide data and compute resources to democratize and fuel best-in-class AI research efforts.
 - o These funds would lay the foundation for broader digital transformation and economic empowerment. Government officials should publicly highlight the impact of this effort and the role of the firms contributing to it.
 - o Similar to the Partnership on AI's work coordinating development of best practices across AI firms,⁸⁵ this effort should be managed by an independent non-profit organization that can link and scale firms' efforts to build digital skills and democratize AI research. The U.S. Digital Service Academy proposed by the Commission could also contribute expertise, volunteers, curriculum development, and other in-kind support.⁸⁶
- **Expand research exchanges between industry and academia.**
 - o Leading technology firms should invest in or expand exchange programs designed to combine top academic talent with world-class private-sector computing resources. Rotational exchanges of this type would both democratize computing access for researchers and simultaneously shed light on new pathways for next-generation AI products that could be commercialized by industry.

Action for the U.S. Bureau of Labor Statistics:

- **Standardize and report data on digital skills in the job market.**
 - o The U.S. Bureau of Labor Statistics should lead an effort in coordination with other agencies such as the Department of Education to collect and regularly update statistics on the digital proficiency of demographic groups and regions, with entries describing specific digital skills needed by firms with job openings. This will enable

academic institutions, firms, and other organizations to prioritize their efforts for educating, reskilling, upskilling, and digital transformation.

Recommendation: Tackle Some of Humanity's Biggest Challenges

Recommendation

If the investments detailed above are implemented, they will set the conditions to harness AI to tackle some of the biggest challenges in science, society, and national security.

Examples of promising initiatives that could improve societal well-being and advance scientific frontiers include, but are not limited to:

- *Enable long-term quality of life.* AI technology that can help the elderly live independently longer, assisting in managing health and daily tasks and improving the quality of life. This can include application of AI to biomedicine to address acute and chronic illnesses and enhance healthy aging.
- *Revolutionize education and lifelong learning.* AI tools that personalize education, training, and retraining at appropriate challenge levels and intuitively evaluate development to optimize standard curricula to promote individual learning success.
- *Transform energy management.* Smart infrastructure for cities that can effectively respond to surges in energy demand and emergencies (both man-made and natural disasters).
- *Effectively predict, model, prepare for, and respond to disasters.* Accurate, near—real time weather, earthquake, and fire line detection and prediction of escalation to aid in emergency response and planning for optimized deployment of limited resources. Autonomous robots for search, rescue, and cleanup in the wake of natural or man-made disaster, providing force-multiplying support to first responders and hazardous materials professionals.

Action for the Office of Science and Technology Policy:

- **Direct the National AI Initiative to align federal investments in AI R&D to tackle significant scientific, technological, and societal challenges.**
 - o The National AI Initiative should identify and oversee realization of opportunities to harness federal R&D investments to take on audacious scientific and technological challenges that could lead to breakthroughs that benefit society and national security.⁸⁷
 - o Prioritization of these efforts should be coordinated with the national security research community and informed by the Technology Competitiveness Council proposed by the Commission⁸⁸ to define areas of research where the application of AI could contribute to progress that provides strategic advantages.

Blueprint for Action: Chapter 11 - Endnotes

¹ *The National Science Foundation: A Brief History*, National Science Foundation (July 15, 1994), <https://www.nsf.gov/about/history/nsf50/nsf8816.jsp#chapter3> (“In fiscal year 1958, the year before Sputnik, the Foundation’s appropriation had leveled at \$40 million. In fiscal 1959, it more than tripled at \$134 million, and by 1968 the Foundation budget stood at nearly \$500 million.”).

² *12 Irreplaceable Innovations Made Possible by NSF*, National Science Foundation (last accessed Feb. 11, 2021), https://www.nsf.gov/news/special_reports/btyb/innovation.jsp. A recent report produced by Computer Science and Telecommunications Board of the National Academies of Sciences, Engineering, and Medicine traces the interplay between fundamental research in information technology (IT) in academia and industry and its effects on capabilities of IT and non-IT sectors. For an illustration of the how the research funded by NSF and others has influenced the technologies that have transformed our everyday lives, see *Information Technology Innovation: Resurgence, Confluence, and Continuing Impact*, National Academies of Sciences, Engineering, and Medicine at 14 (2020), <https://doi.org/10.17226/25961>.

³ We recommend an estimated operating budget of \$20 billion per year. For comparison, NSF has an annual budget of \$8.5 billion (FY 2021), while five U.S. technology firms—Alphabet, IBM, Facebook, Microsoft, and Amazon—spent an estimated \$80.5 billion on AI R&D alone in 2018. See *About the National Science Foundation*, National Science Foundation (last accessed Feb. 11, 2021), <https://www.nsf.gov/about/>; Martijn Rasser, et al., *The American AI Century: A Blueprint for Action*, CNAS (Dec. 17, 2019), <https://www.cnas.org/publications/reports/the-american-ai-century-a-blueprint-for-action>.

⁴ As argued by Donald Stokes in 1997, research should be conceived not as a dichotomy between basic and applied research, but on a quadrant along the axes of “quest for fundamental understanding” and “considerations of use.” Research in the upper-right quadrant is defined as use-inspired basic research—research that advances fundamental knowledge but is driven by a clear purpose. Stokes calls this “Pasteur’s quadrant” after the work of Louis Pasteur, whose research pushed scientific boundaries and had practical applications. See Cherie Winner, *Pasteur’s Quadrant*, Washington State Magazine (2009), <https://magazine.wsu.edu/web-extra/pasteurs-quadrant/>.

⁵ See Chapter 16 of this report for additional details on each of these technologies and why the Commission believes they are critical to future U.S. national competitiveness.

⁶ For additional details on the Commission’s proposed National Technology Strategy and the Technology Competitiveness Council, see Chapter 9 of this report.

⁷ For example, NSF, which provides 85% of federal funding for computer science, funded \$188 million in core AI research in 2019 but did not have room in the budget to fund another \$178 million worth of highly rated proposals. This was an improvement from 2018, when it funded \$165 million but left \$185 million of highly rated work unfunded. Furthermore, NSF (in partnership with the Department of Agriculture) funded seven National AI Research Institutes in 2020 but was unable to fund the more than 30 that were judged worthy of supporting. NSF presentation to NSCAI (January 2020).

⁸ *The Networking & Information Technology Research & Development Program Supplement To The President’s FY2021 Budget*, National Science & Technology Council at 4 (Aug. 14, 2020), <https://www.nitrd.gov/pubs/FY2021-NITRD-Supplement.pdf>.

⁹ *The Networking & Information Technology Research & Development Program Supplement To The President’s FY2020 Budget*, National Science & Technology Council at 11 (Sept. 2019), <https://www.nitrd.gov/pubs/FY2020-NITRD-Supplement.pdf>.

¹⁰ Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

¹¹ The legislation tasks an interagency committee overseen by the National AI Initiative Office to develop every three years a strategic plan for AI that determines and prioritizes areas of AI R&D requiring Federal Government leadership and investment; supports long-term funding for interdisciplinary AI research; provides or facilitates the availability of curated, standardized, secure, representative, aggregate, and privacy-protected data sets for AI R&D; provides or facilitates the necessary computing, networking, and data facilities for AI R&D; supports and coordinates Federal education and workforce training activities; and supports and coordinates the network of AI Research Institutes.

¹² See *Interim Report*, NSCAI at 24-28 (Nov. 2019), <https://www.nscai.gov/previous-reports/>; Craig Willis, *Analysis of Current and Future Computer Science Needs via Advertised Faculty Searches for 2019*, CRA Bulletin (Dec. 7, 2018), <https://cra.org/analysis-of-current-and-future-computer-science-needs-via-advertised-faculty-searches-for-2019/>.

¹³ Expanded funding could go through programs across federal agencies, notably the following. For NSF: CAREER fellowship; Graduate Research Fellowship Program; CyberCorps: Scholarship for Service; Historically Black Colleges and Universities Undergraduate Program; and Research Traineeship. For DoE: Early Career Research Program; Computational Science Graduate Fellowship. For NASA: Space Technology Research Fellowship program. For DoD: DARPA Young Faculty Award; Vannevar Bush Faculty Fellowship; Science, Mathematics, and Research for Transformation Scholarship for Service Program; National Defense Science and Engineering Graduate Fellowship Program; and Historically Black Colleges/Universities and Minority-Serving Institutions Research and Education Program. See Chapter 10's recommendation for the passage of a National Defense Education Act.

¹⁴ Learning techniques such as unsupervised, semi-supervised, self-supervised, one- or zero-shot, and reinforcement learning enable training AI models with less reliance on large data sets of labeled data, albeit often with lower accuracy than with using supervised learning. See Dr. Bruce Draper, *Learning with Less Labeling*, DARPA (last accessed Dec. 19, 2020), <https://www.darpa.mil/program/learning-with-less-labeling>. Reducing reliance on large amounts of labeled data is important when supporting applications where data is scarce or labeling data is cost prohibitive. See *NSCAI Interim Report - Beyond Deep Learning*, NSCAI at 55 (Nov. 2010), https://www.nscai.gov/wp-content/uploads/2021/01/NSCAI-Interim-Report-for-Congress_201911.pdf.

¹⁵ Hybrid AI approaches include integrating statistical machine learning with other techniques such as symbolic AI, knowledge representations, game theory, search, and planning. Hybrid AI approaches are often used in applications of robotics, battle management systems, and resilient systems. NSCAI Staff Correspondence with DARPA (Feb. 22, 2021).

¹⁶ This is a particular challenge for long-lived, autonomous AI systems operating over long durations of time. These systems will likely continuously evolve their mission sets and capabilities, utilizing dynamic learning, along with in-field, in situ updating. All this requires advancing the discipline of TEVV to continuously monitor and ensure such a system's operation remains compliant to performance requirements over its missional lifetime.

¹⁷ The topics were Trustworthy AI, Foundations of Machine Learning, AI-Driven Innovation in Agriculture and the Food System, AI-Augmented Learning, AI for Accelerating Molecular Synthesis and Manufacturing, and AI for Discovery in Physics. The Department of Agriculture teamed with NSF to provide funding toward two of the institutes to support AI research on developing the next generation of and resilience in agriculture. *Artificial Intelligence at NSF*, NSF (Aug. 26, 2020), <https://www.nsf.gov/cise/ai.jsp>.

¹⁸ Around the topics of Human-AI Interaction and Collaboration, Advances in Optimization, AI and Advanced Cyberinfrastructure, Advances in AI and Computer and Network Systems, Dynamic Systems, AI-Augmented Learning, AI to Advance Biology, and AI-Driven Innovation in Agriculture and the Food System. The institutes are funded at a rate of \$4 million per year for five years, totaling \$20 million. See *Id.*

¹⁹ Pub. L. 116-283, sec. 5201(b), William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

²⁰ A 2019 evaluation of the grants made as a component of the National Institutes of Health (NIH) high-risk, high-reward program—which include large, longer-term investments in talent through the NIH Director's Pioneer Award, NIH Director's New Innovator Award, and the NIH Director's Early Independence Award—found that these awards funded highly productive research compared to the work funded under traditional NIH research grants and that they result in a higher technological impact. The high-risk, high-reward program was created to accelerate the pace of biomedical, behavioral, and social science discoveries by supporting creative scientists with highly innovative research. See *Report of the ACD Working Group on High-Risk, High-Reward Research*, National Institutes of Health Advisory Committee to the Director (June 2019), https://www.acd.od.nih.gov/documents/presentations/06132019HRHR_B.pdf.

Blueprint for Action: Chapter 11 - Endnotes

²¹ Studies have found that research that effectively combines diversity of knowledge is more likely to prompt breakthroughs and that interdisciplinary research lends itself to complex problem-solving, developing new research thrusts, and challenging the status quo. See Lee Fleming, *Recombinant Uncertainty in Technological Search*, Management Science (Jan. 2001), <https://funginstitute.berkeley.edu/wp-content/uploads/2012/10/Recombinant-Uncertainty-in-Technological-Search.pdf>; Andrew Barry, et al., *Logics of Interdisciplinarity*, Economy and Society (Feb. 2008), <http://users.sussex.ac.uk/~ir28/IDR/Barry2008.pdf>.

²² The NIH Director's Pioneer Award supports researchers at any career stage who propose bold research projects with unusually broad scientific impact. The program supports awardees with \$3.5 million over five years and requires 51% of time spent on research in the first three years. See *NIH Director's Pioneer Award*, National Institutes of Health (last accessed Jan. 1, 2021), <https://commonfund.nih.gov/pioneer>. Competition for participation in the program is high; reportedly the success rate for applicants is just 1%. See Roberta B. Ness, *The Creativity Crisis*, Oxford University Press at 87 (2015).

²³ Established in 1978, the Howard Hughes Medical Institute (HHMI) supports more than 250 investigators across the United States. Thirty current or former HHMI investigators have been awarded the Nobel Prize. The HHMI Investigator Program is organized around the core belief in the power of individuals to make breakthroughs over time. Through the program, which selects 20 investigators per year, HHMI aims to expand a community of basic researchers and physician scientists who catalyze discovery research in basic and biomedical sciences, plant biology, evolutionary biology, biophysics, chemical biology, biomedical engineering, and computational biology. See *Investigator Program*, HHMI (last accessed Feb. 3, 2021), <https://www.hhmi.org/programs/biomedical-research/investigator-program>; see also *Competition to Select New HHMI Investigators*, HHMI (2020), <https://www.hhmi.org/sites/default/files/programs/investigator/investigator2021-program-announcement-200714.pdf>.

²⁴ This mirrors the HHMI structure and cost model, with HHMI awarding \$8 million over a seven-year term. HHMI updated the length of their award in 2018, extending the term from five to seven years. See *HHMI Bets Big on 19 New Investigators*, HHMI (May 23, 2018), <https://www.hhmi.org/news/hhmi-bets-big-on-19-new-investigators>.

²⁵ Should researchers move institutions over the course of the program, the award would move with them.

²⁶ This could be conducted through a cooperative agreement, mirroring the relationship NSF formed with the Computing Research Association to launch the Computing Innovation Fellows program in 2009 to support postdoctoral PhDs imperiled in finding academic appointments by the downturn of the economy. See *CIFellows*, Computing Community Consortium (last accessed Jan. 1, 2021), <https://cra.org/ccc/leadership-development/cifellows/>. Furthermore, this entity would be able to accept supplemental funding from individuals, corporations, or other non-profits to further strengthen and expand the program.

²⁷ They would provide meaningful feedback to selectees throughout their participation in the program. The quality of feedback provided by reviewers was identified by researchers as a key factor in the success of HHMI investigators. Pierre Azoulay, et al., *Incentives and Creativity: Evidence from the Academic Life Sciences*, NBER (Dec. 2011), <https://www.nber.org/papers/w15466>.

²⁸ Pierre Azoulay & Danielle Li, *Scientific Grant Funding*, MIT & NBER (March 4, 2020), <https://mitsloan.mit.edu/shared/ods/documents/?PublicationDocumentID=6296>. See also the "gold award" model used by the Gates Foundation. *How Grand Challenges Explorations Grants Are Selected*, Bill & Melinda Gates Foundation Global Grand Challenges (last accessed Feb. 3, 2021), <https://gcgh.grandchallenges.org/how-grand-challenges-explorations-grants-are-selected>.

²⁹ The amount of the award would be adjusted in accordance with the specificities of the project. Eligible teams would be composed of researchers based in U.S. academic or research institutions proposing innovative work related to AI.

³⁰ The annual Taulbee Survey that tracks the field of computer science (CS) found that women make up 21.0% of CS bachelor graduates and 20.3% of CS doctoral graduates, and domestic underrepresented minorities account for 14.7% of CS bachelor graduates and only 3.1% of doctoral graduates. Stuart Zweben & Betsy Bizot, *2019 Taulbee Survey*, Computing Research Association at 4-5, 22, (May 2020), <https://cra.org/wp-content/uploads/2020/05/2019-Taulbee-Survey.pdf>.

³¹ Nur Ahmed & Muntasir Wahed, *The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, ArXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>.

³² Joel Klinger, et al., *A Narrowing of AI Research?* ArXiv (Nov. 18, 2020), <https://arxiv.org/pdf/2009.10385.pdf>.

³³ This program may be realized as a single cloud resource or a federation of resources, the pros and cons of which should be considered by the task force with determinations made within their resulting roadmap.

³⁴ Such as the NSF's CloudBank, which brokers cloud access to specific NSF-funded researchers, and the COVID-19 High Performance Computing Consortium, a public-private partnership that grants access to a range of computing resources to serve COVID-19–related research. See CloudBank (last accessed Jan. 2, 2021), <https://www.cloudbank.org/>; the COVID-19 High Performance Computing Consortium (last accessed Jan. 2, 2021), <https://covid19-hpc-consortium.org/>.

³⁵ Pub. L. 116-283, sec. 5106, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

³⁶ For example, AI testbeds could be hosted by DoE's existing national laboratory facilities and high-performance computing resources, by DoD's existing testing and evaluation infrastructure, or by facilities managed by the Department of Transportation, NIH, NIST, or the Department of Agriculture.

³⁷ The National AI Initiative Act of 2020 tasks NIST to develop standards for AI data sharing and documentation. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

³⁸ Through such mechanisms as the Intergovernmental Personnel Act mobility program. *Intergovernmental Personnel Act*, U.S. Office of Personnel Management (last accessed Feb. 1, 2021), <https://www.opm.gov/policy-data-oversight/hiring-information/intergovernment-personnel-act/>.

³⁹ Executive Order 13859 on AI called on federal agencies to “enhance access to high-quality federal data, models, and computing resources to increase their value for AI R&D.” See Donald J. Trump, *Executive Order on Maintaining American Leadership in Artificial Intelligence*, The White House (Feb. 11, 2019), <https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>.

⁴⁰ This would provide for creation of five initial data sets, as well as maintenance over their lifetime and creation of additional data sets as the program matures.

⁴¹ The DoE is well placed to manage such a program, leveraging the cross-disciplinary expertise resident throughout the laboratory network, the unique computing and user facilities housed at the 17 laboratories, and the ability to create and maintain secure data environments. *User Facilities at a Glance*, U.S. Department of Energy: Office of Science (last accessed Jan. 2, 2021), <https://science.osti.gov/User-Facilities/User-Facilities-at-a-Glance#0>. The program could build on the pathfinder Open Data Initiative launched by Lawrence Livermore National Laboratory in partnership with the University of California San Diego, which hosts complex, labelled data sets for testing solutions for scalable ML platforms. See *New Partnerships Results in Increased Access to Compelling “Real World Data,”* UC San Diego (April 21, 2020), <https://library.ucsd.edu/news-events/new-partnership-results-in-increased-access-to-compelling-real-world-data/>; *Open Data Initiative*, Lawrence Livermore National Laboratory (last accessed Jan. 2, 2021), <https://data-science.llnl.gov/open-data-initiative>.

⁴² For example, NSF, NASA, NIH, and DARPA have all sponsored or created data resources relevant to an open knowledge network. In addition, government and community-led efforts to pool data to build solutions to the COVID-19 pandemic could be leveraged.

⁴³ *Open Knowledge Network: Summary of the Big Data IWG Workshop*, National Science & Technology Council (Nov. 2018), <https://www.nitrd.gov/pubs/Open-Knowledge-Network-Workshop-Report-2018.pdf>.

⁴⁴ This would build on ongoing efforts through NSF's Convergence Accelerator track on Open Knowledge Networks. *NSF Convergence Accelerator Awards Bring Together Scientists, Businesses, Nonprofits to Benefit Workers*, NSF (Sept. 10, 2019), https://www.nsf.gov/news/special_reports/announcements/091019.jsp.

Blueprint for Action: Chapter 11 - Endnotes

⁴⁵ For additional details and recommendations on technologies associated with AI that are important to U.S. technology leadership, see Chapter 16 of this report. A strategic industry is considered by the government to be very important to a country's economy or safety. In the national security context, it is considered critical to the country's competitive advantage over an adversary. While the United States' 16 critical infrastructure sectors refer to large segments of the economy "whose assets, systems, and networks, whether physical or virtual, are considered so vital to the United States," a strategic industry refers to a much more specific group of companies or businesses. See *Critical Infrastructure Sectors*, Cybersecurity and Infrastructure Agency (last accessed Jan. 4, 2020), <https://www.cisa.gov/critical-infrastructure-sectors>; see also *Strategic Industry*, Cambridge Dictionary (last accessed Jan. 4, 2020), <https://dictionary.cambridge.org/dictionary/english/strategic-industry>.

⁴⁶ U.S. federal agencies collectively have an annual information technology (IT) budget of \$90 billion—one-tenth the annual revenue of the top five U.S. tech firms—yet the majority of government systems are "outdated and poorly protected." *An American Budget*, U.S. Office of Management and Budget at 9 (Feb. 2018), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2018/02/budget-fy2019.pdf>.

⁴⁷ Congress, in the Consolidated Appropriations Act, 2021, called on the General Services Administration (GSA) to create a five-year program to be known as the "AI Center of Excellence" to "(1) facilitate the adoption of artificial intelligence technologies in the Federal Government; (2) improve cohesion and competency in the adoption and use of artificial intelligence within the Federal Government; and (3) carry out paragraphs (1) and (2) for the purposes of benefiting the public and enhancing the productivity and efficiency of Federal Government operations." *Rules Committee Print 116-68, Text of the House Amendment to Senate Amendment to H.R. 133*, U.S. House Committee on Rules at 378 (Dec. 21, 2020), <https://rules.house.gov/sites/democrats.rules.house.gov/files/BILLS-116HR133SA-RCP-116-68.pdf> (referring specifically to section 103 of the Consolidated Appropriations Act, 2021).

⁴⁸ The Consolidated Appropriations Act, 2021, outlines AI CoE's duties to include "advising the Director of the Office of Science and Technology Policy on developing policy related to research and national investment in artificial intelligence." *Rules Committee Print 116-68, Text of the House Amendment to Senate Amendment to H.R. 133*, U.S. House Committee on Rules at 380 (Dec. 21, 2020), <https://rules.house.gov/sites/democrats.rules.house.gov/files/BILLS-116HR133SA-RCP-116-68.pdf>.

⁴⁹ The National AI Initiative Act of 2020 directs the Director of OSTP to establish a "National Artificial Intelligence Initiative Office" within OSTP to "(1) provide technical and administrative support to the Interagency Committee and the Advisory Committee; (2) serve as the point of contact on Federal artificial intelligence activities carried out under the Initiative for Federal departments and agencies, industry, academia, nonprofit organizations, professional societies, State governments, and such other persons as the Initiative Office considers appropriate to exchange technical and programmatic information; (3) conduct regular public outreach to diverse stakeholders, including through the convening of conferences and educational events, the publication of information about significant Initiative activities on a publicly available website, and the dissemination of findings and recommendations of the Advisory Committee, as appropriate; and (4) promote access to and early adoption of the technologies, innovations, lessons learned, and expertise derived from Initiative activities to agency missions and systems across the Federal Government, and to industry, including startup companies." Pub. L. 116-283, sec. 5102, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁵⁰ The Consolidated Appropriations Act, 2021, outlines AI CoE's duties to include "advising the Administrator, the Director, and agencies on the acquisition and use of artificial intelligence through technical insight and expertise, as needed." *Rules Committee Print 116-68, Text of the House Amendment to Senate Amendment to H.R. 133*, U.S. House Committee on Rules at 379 (Dec. 11, 2020), <https://rules.house.gov/sites/democrats.rules.house.gov/files/BILLS-116HR133SA-RCP-116-68.pdf>.

⁵¹ The Consolidated Appropriations Act, 2021, outlines AI CoE's duties to include "(1) regularly convening individuals from agencies, industry, Federal laboratories, nonprofit organizations, institutions of higher education, and other entities to discuss recent developments in artificial intelligence, including the dissemination of information regarding programs, pilots, and other initiatives at agencies, as well as recent trends and relevant information on the understanding, adoption, and use of artificial intelligence; (2) collecting, aggregating, and publishing on a publicly available website information regarding programs, pilots, and other initiatives led by other agencies and any other information determined appropriate by the Administrator." *Rules Committee Print 116-68, Text of the House Amendment to Senate Amendment to H.R. 133*, U.S. House Committee on Rules at 378-79 (Dec. 21, 2020), <https://rules.house.gov/sites/democrats.rules.house.gov/files/BILLS-116HR133SA-RCP-116-68.pdf>.

⁵² The digital playbook should consider recent upgrades made to Acquisition.gov. An important element of the development of an effective playbook is ensuring its interface and content accounts for different user profiles. Critical among those user profiles is that of a small-business or non-traditional government contractor that may be unfamiliar with the process to even begin eligibility for a government contract. *Access the Federal Acquisition Regulation*, U.S. General Services Administration (last accessed Feb. 18, 2021), <https://www.acquisition.gov/>.

⁵³ The SBIR program is one of the largest and longest-standing programs for federally funded R&D in small businesses. It was established in 1982 as part of the Small Business Innovation Development Act, and Federal agencies with extramural research and development budgets that exceed \$100 million set aside 3.2% of their budgets to fund the SBIR program. The program is structured in three phases: Phase I awards of approximately \$50,000 to \$250,000 for six months to vet "technical merit, feasibility, and commercial potential"; Phase II awards of \$750,000 to \$1,700,000 for two years to support successful efforts initiated in Phase I; and Phase III, which is not funded by SBIR dollars, to pursue commercialization. The program issues a higher number of Phase I awards but allocates more funding toward Phase II, with the goal of placing many small bets on novel technologies and only scaling those that show real promise. NSCAI Engagement (Sept. 25, 2020); see also *About, Small Business Innovation Research* (last accessed Feb. 3, 2021), <https://www.sbir.gov/about>.

⁵⁴ For example, AFWERX's Supplemental Funding Pilot Program (TACFI and STRATFI) and USD(R&E)'s Accelerated Transition funding program.

⁵⁵ "As of November 2020, agencies may issue a Phase I award (including modifications) up to \$259,613 and a Phase II award (including modifications) up to \$1,730,751 without seeking SBA approval. Any award above those levels will require a waiver." *About, Small Business Innovation Research* (last accessed Feb. 3, 2021), <https://www.sbir.gov/about>. See also *Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Program Policy Directive*, U.S. Small Business Association (May 2, 2019), https://www.sbir.gov/sites/default/files/SBIR-STTR_Policy_Directive_2019.pdf.

⁵⁶ The Air Force, in partnership with Air Force Research Lab (AFRL) and the National Security Innovation Network (NSIN), developed Open SBIR Topics, which includes a "few big bets" (strategic financing): rewards of up to \$15 million, with 1:1:2 Program-SBIR-Private Matching options. *SBIR Open Topics*, U.S. Air Force AFWERX (last accessed Feb. 3, 2021), <https://www.afwerx.af.mil/sbir.html>.

⁵⁷ Specifically, on page 74 of the SBA SBIR/STTR Policy Directive, the line "For example, some agencies administer Phase IIB awards that differ from the base Phase II in that they require third party matching of the SBIR/STTR funds." could be changed to "For example, some agencies administer Phase II or IIB awards that require third party matching of the SBIR/STTR funds." *Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) Program Policy Directive*, U.S. Small Business Administration at 74 (May 2, 2019), https://www.sbir.gov/sites/default/files/SBIR-STTR_Policy_Directive_2019.pdf.

⁵⁸ The current application portals for beta.sam.gov and the "Defense SBIR/STTR Innovation Portal" are significant barriers to entry for non-traditionals trying to work with the DoD. NSCAI staff engagement (Feb. 9, 2021).

⁵⁹ *Future of Defense Task Force Report 2020*, U.S. House Committee on Armed Services at 68 (Sept. 23, 2020), https://armedservices.house.gov/_cache/files/2/6/26129500-d208-47ba-a9f7-25a8f82828b0/424EB2008281A3C79BA8C7EA71890AE9.future-of-defense-task-force-report.pdf.

⁶⁰ For more examples of innovation initiatives and recommendations to scale their impact, see Chapter 2 of this report and its associated Blueprint for Action.

Blueprint for Action: Chapter 11 - Endnotes

⁶¹ Contracts must be easier to understand and fill out, review periods shortened and clearly communicated, and oversight streamlined to keep pace with the current rate of technology innovation.

⁶² Phase II and supplemental awards should be based on a broader diligence process that includes the long-term health and viability of the company. This assessment should consider as a starting point the firm's technical capabilities, financial structure, management structure, and the larger commercial market opportunities.

⁶³ Phase II enhancements, sometimes called Phase IIB/II.5 contacts, have become a common method to extend SBIR dollars to promising projects that fail to secure Phase III funding. The Navy Commercialization Readiness Program oversees the distribution of Phase II.5 contracts "to further develop SBIR technologies and to accelerate transition for existing Phase II projects." *Navy Phase II.5 Structure and CRP*, U.S. Navy (last accessed Feb. 3, 2021), <https://www.navysbir.com/cpp.htm>. The Air Force's AFWERX, Army, and DARPA, as well as several Federal agencies outside the DoD, also use Phase IIB awards. The Office of the Secretary of Defense (OSD) Transitions SBIR Technology Pilot Program provides SBIR awardees the opportunity to apply for Phase II Enhancement (e) and Accelerated Transition funding for the funding sponsor. However, current funding limits set by SBA reduce their efficacy by including Phase II enhancements under the Phase II cap of SBIR dollars. NSCAI staff engagement (Sept. 23, 2020). For further detail, see *Interim Report and Third Quarter Recommendations*, NSCAI at 52-57 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

⁶⁴ This effort would be informed by the Technology Annex to the National Defense Strategy recommended in Chapter 2 of this report.

⁶⁵ The future will likely be defined by a fusion of many different AI approaches including expert systems, model-based AI, symbolic-based AI, statistical ML, and new and evolving AI approaches such as neurosymbolic AI. See *Neuro-Symbolic AI*, MIT-IBM Watson AI Lab (last accessed Feb. 3, 2020), <https://mitibmwatsonailab.mit.edu/category/neuro-symbolic-ai/>.

⁶⁶ DARPA's SBIR program, for example, is unique in its long time horizon. Most of its investments are pre-commercial and will take another eight to 10 years to develop before results can be scaled for military or commercial use.

⁶⁷ David E. Cooper, *Defense Industry Consolidation: Competition Effects of Mergers and Acquisitions*, Statement before the U.S. Senate Committee on Armed Services Subcommittee on Acquisition and Technology (March 4, 1998), <https://www.gao.gov/assets/110/107240.pdf>.

⁶⁸ For example, Intel's recent chip missteps have jeopardized U.S. leadership in the design and manufacturing of advanced semiconductors. See Michael Kan, *Intel: Sorry, But Our 7nm Chips Will Be Delayed to 2022, 2023*, (July 23, 2020), <https://www.pcmag.com/news/intel-sorry-but-our-7nm-chips-will-be-delayed-to-2022-2023>.

⁶⁹ For example, small firms have difficulty affording the cost of compute resources and data for training sophisticated ML models. Nur Ahmed & Muntasir Wahed, *The De-democratization of AI: Deep Learning and the Compute Divide in Artificial Intelligence Research*, arXiv (Oct. 22, 2020), <https://arxiv.org/abs/2010.15581>.

⁷⁰ Michael Porter, *The Competitive Advantage of Nations*, Harvard Business Review (1990), <https://hbr.org/1990/03/the-competitive-advantage-of-nations>.

⁷¹ *Interim Report*, NSCAI at 26 (Nov. 2019), <https://www.nscai.gov/previous-reports/>; see also Michael Porter, *Clusters and the New Economies of Competition*, Harvard Business Review (1998), <https://hbr.org/1998/11/clusters-and-the-new-economics-of-competition>.

⁷² William R. Kerr & Frederic Robert-Nicoud, *Tech Clusters*, Journal of Economic Perspectives at 63 (2020), <https://pubs.aeaweb.org/doi/pdfplus/10.1257/jep.34.3.50>.

⁷³ The program office could be modeled on the Advanced Manufacturing National Program Office that coordinates Manufacturing USA, a network of manufacturing innovation institutes. See Manufacturing USA (last accessed Feb. 3, 2021), <https://www.manufacturingusa.com/>.

⁷⁴ For example, proximity to research facilities operated by the departments of Defense and Energy or access to technically oriented military installations should be prioritized.

⁷⁵ See the Chapter 13 Blueprint for Action for more details on the importance of U.S. access to trusted and assured microelectronics for national security use cases. The Commission also proposes a preliminary list of strategic emerging technologies that are critical to U.S. national competitiveness in Chapter 16 of this report.

⁷⁶ The existence of a nascent cluster suggests industry has already passed the market test. Mark Muro & Bruce Katz, *The New “Cluster Moment”: How Regional Innovation Clusters Can Foster the Next Economy*, The Brookings Institution (Sept. 21, 2010), <https://www.brookings.edu/research/the-new-cluster-moment-how-regional-innovation-clusters-can-foster-the-next-economy/>. Resources like the U.S. Cluster Mapping Project will also be essential to identify which locations are economically viable. See U.S. Cluster Mapping (last accessed Feb. 3, 2021), <http://clustermapping.us/>.

⁷⁷ Anchor institutions are firms, not-for-profit institutions, and research universities that locate near the cluster and pursue joint R&D with federal agencies or other cluster participants.

⁷⁸ *Overview: The New Federal Role in Innovation Clusters, Clustering for 21st Century Prosperity: Summary of a Symposium*, The National Academies Press (2012), <https://www.nap.edu/read/13249/chapter/3#31>.

⁷⁹ For example, the clusters may be co-located with DoE’s national laboratories or military test ranges.

⁸⁰ According to the Council of Economic Advisors, Opportunity Zones (OZs) incentivize private investment in low-income communities by lowering capital gains taxes on businesses investing in the region, which could be a revenue-neutral way of lifting people out of poverty due to the expected reduction in transfer payments. Investors receive tax benefits for investing in Qualified Opportunity Funds, which can be used to make equity investments in partnerships or corporations that operate in an OZ. The funds can also be used to purchase tangible property for use in the fund’s trade or business. See *The Impact of Opportunity Zones*, The Council of Economic Advisors (Aug. 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/08/The-Impact-of-Opportunity-Zones-An-Initial-Assessment.pdf>.

⁸¹ Private-sector contributions may comprise cost sharing in joint R&D projects, donations, or membership dues, if such a model is adopted.

⁸² CODE Advocacy Coalition (last accessed Jan. 2, 2021), <https://advocacy.code.org/>.

⁸³ See, e.g., *Microsoft Philanthropies: TechSpark*, <https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE4s6AL>; Carolina Milanesi, *STEM Education as a Diversity Driver in Tech*, Amazon (Sept. 14, 2020), <https://www.aboutamazon.com/news/community/stem-education-as-a-diversity-driver-in-tech>; *Applied Digital Skills: Teach and Learn Practical Digital Skills*, Google (last accessed Jan. 2, 2021), <https://applieddigitalskills.withgoogle.com/s/en/home>.

⁸⁴ Michael Wade, *Corporate Responsibility in the Digital Era*, MIT Sloan Management Review (April 28, 2020), <https://sloanreview.mit.edu/article/corporate-responsibility-in-the-digital-era/>.

⁸⁵ Partnership on AI has a mission to shape best practices, research, and public dialogue about AI’s benefits for people and society, with partners from more than 100 companies and research organizations. *Partnership on AI* (last accessed Feb. 3, 2021), <https://www.partnershiponai.org>.

⁸⁶ See Chapter 6 of this report for further discussion of the Commission’s proposed U.S. Digital Service Academy.

⁸⁷ One way this could be enacted is by assigning “national mission managers” to oversee each opportunity identified.

⁸⁸ As recommended in Chapter 9 of this report.

Chapter 12: Intellectual Property

Blueprint for Action

America's intellectual property (IP) laws and institutions must be considered as critical components for safeguarding U.S. national security interests, including advancing economic prosperity and technology competitiveness. Prioritization of IP policy is especially important given China is both leveraging and exploiting IP policies as a tool within its national strategies for emerging technologies. The United States must, at a minimum, articulate and develop national IP reforms and policies with the goal of incentivizing, expanding, and protecting artificial intelligence (AI) and emerging technologies,¹ at home and abroad. Such policies should be developed and proposed via the Executive Branch with a process that integrates the disparate departments and agencies that serve important roles in promoting U.S. innovation.

Recommendation

Recommendation: Develop and implement national IP policies and regimes to incentivize, expand, and protect AI and emerging technologies as part of national security strategies.

Action for the President:

- **Issue an Executive Order to prioritize IP policies for AI and critical emerging technologies.**
 - o The President should issue an Executive Order to recognize IP policy as a national priority and establish a comprehensive process to reform and establish new IP policies and regimes for AI and critical emerging technologies that further national security, economic, and technology competitiveness strategies.
 - o The Executive Order should:
 - Direct the Vice President, as Chair of the Technology Competitiveness Council (TCC)² or otherwise as chair of an interagency task force,³ to oversee the comprehensive process;
 - Direct the Secretary of Commerce to:
 - Lead, on an ongoing basis, the development of proposals (Executive and/or Legislative Branch actions) to reform and establish new IP policies and regimes to incentivize, expand, and protect AI and emerging technologies;
 - In executing these responsibilities, coordinate with the Under Secretary of Commerce for Intellectual Property, the Director of the U.S. Patent and Trademark Office (USPTO), and other relevant Executive Branch agencies;

consult with the Director of the U.S. Copyright Office; and convene public deliberations, to include at a minimum academia and industry;

- Direct the USPTO Director, in his capacity as advisor to the President,⁴ to:
 - Submit, within 90 days, a report to the Vice President, in their capacity as the head of the TCC or interagency task force, that (1) identifies and analyzes metrics, trends, and data necessary to inform IP policymaking, particularly as prioritized in the Executive Order; and (2) identifies the associated U.S. Executive Branch departments and agencies that will be required to provide any requisite data;
 - Submit, within 12 months from issuance of the first report, a second report, or portions on a rolling basis, to the Vice President that (1) comprehensively assesses the weaknesses in the current U.S. IP policies and regimes, relative to IP regimes of other nations, for incentivizing, expanding, and protecting innovation in AI and emerging technologies and supporting national strategies; (2) examines the non-exhaustive list of “IP considerations” (see second recommendation); and (3) proposes corresponding executive and legislative actions for reforming and establishing new IP policies and regimes;
 - Provide all necessary information and advice to the Vice President to enable a fulsome analysis of the IP proposals;
- Direct the Vice President to:
 - Lead an ongoing assessment of IP policies, regimes, and reform proposals from the Secretary of Commerce that should be implemented and integrated into national security, economic, and technology competitiveness strategies;
 - Empower the Secretary of Commerce to facilitate implementation of IP policies and regimes assessed as critical to national security, economic, and technology competitiveness strategies; and
- Direct Executive Branch departments and agencies to resource and support the Secretary of Commerce in executing these Executive Order efforts, including providing the identified metrics and trends.

Actions for the Secretary of Commerce and USPTO Director:

- **Establish, as necessary, in consultation with the Director of the USPTO, a committee of multidisciplinary experts, from within and outside the U.S. government, to provide technical and IP-related expertise and advice in implementing this Executive Order.**
- **Convene public deliberations, to include at a minimum academia and industry, in executing these Executive Order responsibilities. The outcome of these deliberations should inform proposed IP policies and regimes.**
- **Assess metrics and data necessary to inform IP policy.**
 - In assessing the proper metrics and data necessary to inform IP policy deliberation as required by the Executive Order, the Secretary of Commerce and USPTO

Director should take a whole-of-government approach. Due to the breadth of the IP considerations, including those delineated in this report, as well as the far-reaching impact of IP upon many segments of the U.S. economy and innovation ecosystem, there are many U.S. government entities that may already track relevant metrics or have the capability to expand their analyses to address the necessary prioritization of IP for AI and emerging technologies.

- For example, innovation and investment trends based on patent filings, and, where possible, licensing data—in various technology sectors, including by foreign countries, particularly China—should be analyzed (e.g., to assess quality and research trends⁵), with care not to rely solely on patent counting.
- Other potential metrics include but are not limited to tracking of patents self-declared as standard essential in comparison to patents actually licensed; licensing to unrelated parties; the impact of prior art on the U.S. patent and trademark examination systems; international filings for IP protections on U.S.-funded research, particularly without U.S. funders' or inventors' awareness; the ratio of U.S. companies filing for IP protections, as well as pursuing IP-related litigation, in the U.S. versus abroad; and patent assignment data.

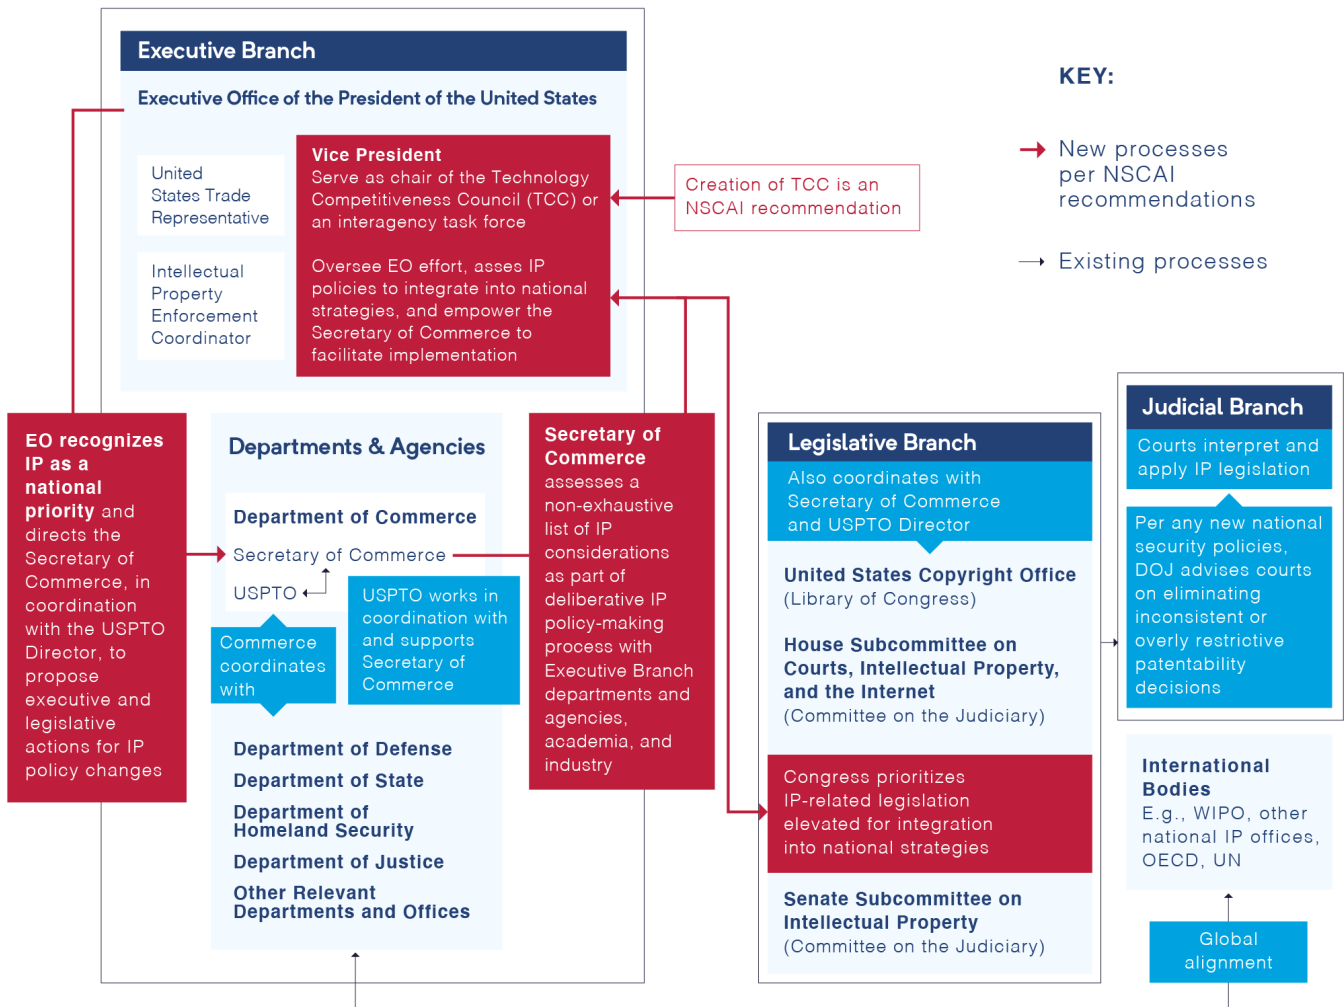
Action for the Department of Justice:

- **Advise courts on ensuring consistency on patentability decisions.**
 - The Department of Justice, through the Solicitor General and the Civil Appellate Section, should advise federal courts on eliminating confusing, inconsistent, or overly restrictive patentability decisions to ensure consistency with national security policies.

Action for Congress:

- **Prioritize proposed IP-related legislation to bolster U.S. national strategies, including for national security, economic interests, and technology competitiveness.**
 - Congress should prioritize legislative recommendations for IP policies and regimes elevated by the Vice President, as Chair of the TCC or an interagency task force. This is particularly important given Congress is responsible for passing patent and IP legislation that the USPTO and other relevant stakeholders execute and follow. Additionally, the U.S. Copyright Office is housed as a federal department within the Library of Congress as the principal advisor to Congress on copyright matters and administers copyright registrations.⁶

Executive Order to Prioritize IP Policies for AI and Emerging Technologies*



*This illustration is not comprehensive of all relevant U.S. government entities with intellectual property responsibilities

Recommendation

Recommendation: The Secretary of Commerce should assess and examine the following non-exhaustive list of “IP considerations,” in coordination with the Under Secretary of Commerce for IP and the Director of the USPTO, as part of developing and proposing reforms and new IP policies and regimes to the Vice President.

Action for the Secretary of Commerce:

- **Assess and examine the following non-exhaustive list of 10 considerations for intellectual property as part of the reports submitted to the Vice President as mandated by the Executive Order.**

1. *Patent Eligibility:* The Secretary of Commerce should assess and articulate the impact of current patent eligibility laws on innovation in AI and emerging technologies from an economic, trade, and national security policy perspective to better inform the legislative and agency efforts on patent eligibility reform. America’s IP regime has spurred American ingenuity since the late 18th century. By protecting “any new and useful process, machine, manufacture, or composition of matter” through stable legal institutions governed by the rule of law, inventors and investors have relied on America’s IP system to provide the certainty necessary to justify large and risky R&D investments,⁷ which are critical for technologies.⁸ A strong and robust patent system is equally critical to incentivizing American innovation in AI and emerging technologies that affect national security.⁹ Unfortunately, recent patent eligibility court rulings have narrowed the scope of inventions that are eligible for patent protection. This has resulted in a broad swath of innovation that is now ineligible for patent protection in both digital technologies and biopharma, among others.¹⁰ The legal uncertainty for U.S. innovators and companies as to whether their inventions will be eligible for patent protection or susceptible to invalidation once granted is pervasive.¹¹ This uncertainty in turn has impacted investments in AI and technologies critical to national security. Empirical studies have proven that patents are causally linked to venture capital investments in startups, and, as a result, are causally linked to the success of startups.¹² Recent reports, however, reveal that investments in patent-intensive U.S. startups that develop critical technologies (e.g., computer hardware, semiconductors, medical devices and supplies, and pharmaceuticals and biotechnology) have declined relative to non-patent-intensive companies.¹³ This is consistent with investors consistently reporting that patent eligibility is a key factor in their decisions whether to invest in a particular company’s technologies or bring a new product to market.¹⁴

Legislation appears to be the only practical means to reform patent eligibility doctrine. The Judiciary, specifically the Supreme Court, has indicated an unwillingness to revisit its decisions in the past decade that have created this fundamental problem in patent eligibility doctrine.¹⁵ The USPTO has adopted a framework for assessing patent eligibility during the examination process of patent applications, which has had positive results in providing greater certainty to patent applicants,¹⁶ but the Federal Circuit does not seem inclined to follow USPTO guidance.¹⁷

Efforts to reform the patent eligibility doctrine by amending the relevant provision in the patent statutes failed in 2019.¹⁸ Efforts continue to restart the legislative reform process. A national security point of view has not been expressed on the impact of patent eligibility law on technologies critical to national security, such as AI, microelectronics, 5G telecommunications, quantum computing, and biotechnology. A national security point of view on the impact of current patent eligibility laws on AI and emerging technologies should inform a national IP strategy.

2. Counter China's narrative on winning the innovation competition: The Secretary of Commerce, in coordination with relevant departments and agencies (e.g., Department of State, USTR), should address how the United States might best counter China's efforts to shape the narrative that it is winning the innovation competition based in part on its patent application filings and other interventions in its technology markets.¹⁹ China has become the domestic forum with the highest number of patent application filings, and China's companies and inventors are the most prolific AI patent application filers globally.²⁰ This benchmark helps to shape the narrative that China has become the leader in innovation because intensive patenting has been shown to generally correlate to economic growth.²¹ China also is garnering this reputation when it comes to emerging technologies such as AI.²² Sources claim that China is outpacing the United States in filing worldwide AI-related patent applications.²³ However, high levels of patenting output is not necessarily indicative of high levels of inventive output.²⁴ Specifically, non-market factors driven by state-sponsored interferences can distort filings.²⁵ Moreover, China often files patents as a "numbers game," which can lead to mischaracterizing its technological prowess. Similarly, China's 5G companies declare the most patents as "standard essential," appearing to marry China's concerted, top-down strategy to advance its AI and emerging technology agenda by influencing international standards setting with its goals to dominate numeric benchmarks.²⁶ The Secretary of Commerce should examine what measures need to be undertaken to counterbalance the narrative of China's technological dominance based on selective patenting data.

3. Impact of China's patent application filings on USPTO and U.S. inventors: The Secretary of Commerce, in coordination with the USPTO Director, should assess whether the USPTO requires additional resources, both human and technical, to ensure high-quality patent examination and recommend policies to address any concerns. In doing so, the Secretary of Commerce should assess the impacts of increased filings from China and AI-generated prior art (the term in patent law for the worldwide scientific and technical knowledge by which an invention is evaluated to determine if it is new). The large body of often low-quality prior art created by China's high-volume patenting has the potential to adversely impact global patent examination systems, including those of the USPTO.²⁷ At the same time, U.S. inventors may face hurdles in patenting around massive amounts of low-quality Chinese prior art.²⁸ The USPTO has also noted that stakeholders have raised the issues of whether AI may generate a proliferation of prior art, making it difficult to find relevant prior art for examination.²⁹

4. *Impediments to AI public-private partnerships and international collaboration:* The Secretary of Commerce should assess any impediments to the IP contractual ecosystem to strengthen AI partnerships among national security departments and agencies, industry, and international collaboration. This should include assessing and addressing ambiguities in the Federal Acquisition Regulation and the Defense Federal Acquisition Regulation Supplement relevant to AI and data. AI development presents unique IP contractual issues. For example, industry AI developers will likely need access to relevant U.S. Government training data to develop AI-enabled government solutions or applications. If the solution or application is dual-use, the private entity may want to provide a license for the U.S. Government agency to access the AI application, but retain the IP in the AI model to license to others. But there are unanswered questions as to whether the U.S. Government agency has any IP rights or ownership in the model that was trained on its data.³⁰ The U.S. Government agency may also want to retain IP rights in order to avoid “vendor lock.”³¹ These outstanding questions about IP rights and ownership issues could also arise in international AI system R&D collaboration, where impediments can be amplified by conflicting national laws on IP and/or data protections.

5. *IP protection for data:* The Secretary of Commerce should assess whether there is a need for *sui generis* protection or additional IP-type of protections for data and propose policies and/or legislation if protection is deemed necessary. Data is critical to AI and machine learning (ML), but gaps may exist in current protection regimes afforded by patent or copyright. Inadequate protections for data may disincentivize the necessary investments in developing these critical data sets as well as public disclosure and sharing agreements.³² While protections for data might be a future need, the U.S. should be proactive in assessing and addressing the necessity of such protections. The Secretary of Commerce also should explore ways to protect and incentivize creation of data sets while allowing the data to be shared at some point, particularly with smaller entities that might not otherwise be able to enter the market.³³ An analysis of the strengths and weaknesses of the European *sui generis* database protections should inform this assessment.³⁴

6. *Combat IP theft:* The Secretary of Commerce, in coordination with relevant departments and agencies (e.g., USTR, Intellectual Property Enforcement Coordinator, the National Science Foundation, the Office of Science and Technology Policy,³⁵ as well as the Departments of Homeland Security,³⁶ Justice,³⁷ and State) should assess and identify additional efforts that the Executive Branch should undertake to counter IP theft threats, including actions in collaboration with allies and partners.³⁸ In particular, the Secretary of Commerce should clearly articulate that the U.S. counter-IP theft strategy will contain both criminal and civil economic dimensions. The Department of Commerce should utilize all available tools for establishing a deterrence regime to punish firms guilty of stealing U.S. IP and deter future IP theft to level the playing field for U.S. and allied firms. These tools should include placing offending companies on the Bureau of Industry & Security entity list,³⁹ blocking visas of key employees, or levying tariffs against products derived from stolen IP. Solutions that should be explored include training for allies and partners to stop

counterfeits at borders and efforts to increase individuals' respect for IP and recognition of and ways to avoid counterfeits. In addition, the Secretary should assess methods and means for strengthening and updating existing mechanisms available to American victims of trade-secret theft, including reintroducing legislation to strengthen the Economic Espionage Act by, for example, increasing damages available to trade-theft victims and extending the statute of limitations.⁴⁰

7. Inventorship by AI: The Secretary of Commerce should assess the need for policy changes for issues raised by AI-generated inventions and creations, particularly as technologies evolve. The USPTO has determined that under current legal doctrine, an inventor must be a natural person and denied a patent application naming a machine as the inventor.⁴¹ The U.S. is not alone in this position.⁴² The USPTO also issued extensive requests for public comments on a variety of AI IP policy issues, including AI's impact on inventorship and ownership, as well as impacts on non-patent IP protections, such as copyright. As a result, the USPTO issued a comprehensive report of public views on AI and IP policy. The majority of commenters agreed that, given that current AI capabilities are limited to "narrow AI" (AI systems that are trained and perform individual tasks in well-defined domains) and artificial general intelligence is not yet a reality, current AI could neither invent nor author without human intervention.⁴³ The Secretary of Commerce should consult with allies and partners to ensure continued harmonization around the various IP issues raised by AI-generated inventions and creations and gain an understanding of China's strategies for addressing these issues, particularly as AI technologies move past narrow AI.

8. Global IP alignment: The Secretary of Commerce, in coordination with relevant departments and agencies (e.g., USPTO, IPEC, USTR, Department of Defense, Department of State), should work with partners and allies to develop global disincentives for IP theft and alleviate any inconsistencies in patent regimes that make it overly difficult for companies to protect their patents in multinational markets. In doing so, the Secretaries should leverage the Commission's recommendation that the United States and allies—through the Emerging Technology Coalition—explore coordinated approaches to IP (as part of the NSCAI-proposed critical area No. 4: Promoting and Protecting Innovation⁴⁴), including a mutual agenda within the WIPO's Conversation on AI and IP and forums with broader mandates. The Secretaries also should assess whether current forums for dialogues on global IP alignment are sufficient or whether new forums or venues are necessitated, particularly given any changes to domestic IP policies or regimes identified during the review of the other IP considerations. For example, if the U.S. determines new protections or policies are needed for data, it may need to work with key allies and partners—bilaterally and multilaterally—to ensure global harmonization.

9. Democratize innovation and IP ecosystems: The Secretary of Commerce should assess whether additional Executive Branch efforts are necessary to expand the innovation base and democratize access to and create more jobs in the innovation and IP ecosystem.⁴⁵

The USPTO, in collaboration with the Secretary of Commerce, has undertaken initiatives to expand the U.S. innovation base by creating the National Council for Expanding American Innovation (NCEAI) to develop a comprehensive national strategy to increase equity and fuel the U.S. innovation ecosystem by encouraging, empowering, and supporting all future innovators.⁴⁶ The Secretary of Commerce should ensure that the USPTO has the full support of the Executive Branch in these initiatives. As part of the NCEAI initiative, the Secretary of Commerce also should focus on assessing and identifying potential actions and tools that can fast-track processes and streamline guidance for startups seeking IP protections and ensuring resources for assisting small and medium-sized entities. Such a focus is particularly important when comparing the impact of litigation costs and potentially overly burdensome processes in the U.S., relative to other countries, on U.S. inventors' decisions to pursue IP protections in the United States.⁴⁷

10. *“Standard essential” patents process*⁴⁸: The Secretary of Commerce, in coordination with relevant departments and agencies (e.g., USPTO, NIST, and the Department of State), should assess policies by which the U.S. can serve a leadership role in and ensure U.S. firms are able to fully participate in the processes by which “standard essential” patents are claimed and asserted.⁴⁹ This would help ensure the continuing legitimacy of the standard-setting process, a privately developed method for efficiently coordinating development

and deployment of new technologies in the marketplace, and deflect Beijing's attempt to dominate or manipulate these processes through its own coordination of firms from China. Chinese Communist Party leadership has articulated a linkage between patent leadership in emerging technologies like AI and the standards-setting processes for these same technologies.⁵⁰ Current trends confirm China's intention to use both patents and standards to lead in technological innovation.⁵¹ Additional mechanisms may be necessary to protect the integrity of international standards-setting as well as to protect and promote U.S. innovation, such as identifying efforts by foreign governments to influence, directly or indirectly, standard-setting organizations. This would also include identifying foreign governments subsidizing or otherwise incentivizing the over-declaration of patents as "standard essential"⁵² or creating barriers to U.S. participation in foreign standard-setting bodies. The Secretary of Commerce also should explore how the U.S. government might support smaller U.S. companies and inventors fully participating in the standard-setting process and encourage the observation of licensing or legal disputes in foreign jurisdictions by U.S. government officials from U.S. Embassies and Missions. Relatedly, the Secretary of Commerce, in coordination with the Director of the USPTO, should assess foreign court rulings on licensing that may impact U.S. national sovereignty to determine a coherent U.S. position or response.⁵³

Blueprint for Action: Chapter 12 - Endnotes

¹ For a discussion of the U.S. government's efforts to define and prioritize critical emerging technologies, as well as the Commission's recommended eight emerging technologies key to U.S. national competitiveness, see Chapter 16 of this report and its associated Blueprint for Action.

² NSCAI recommended the creation of a Technology Competitiveness Council in its 2020 Interim Report and Third Quarter Recommendations. See *Interim Report and Third Quarter Recommendations*, NSCAI at 180 (Oct. 2020), <https://www.nscai.gov/previous-reports/> ("Technology Competitiveness Council, led by the Vice President and with a Commissioned Assistant to the President as the day-to-day coordinator, to fill this role.") If the TCC is not established as recommended by the Commission, the Commission recommends that the Vice President should lead these efforts.

³ If the TCC is not established, the President, through an Executive Order, should establish a task force to address the mandate recommended here.

⁴ The USPTO Director "shall advise the President, through the Secretary of Commerce, on national and certain international intellectual property policy issues." 35 U.S.C. § 2.

⁵ As an example, an examination of China's patents can provide insight into its biotechnology and genomics research and plans. See Kristy Needham, *Exclusive: China Gene Firm Providing Worldwide COVID Tests Worked with Chinese Military*, Reuters (Jan. 30, 2021), <https://www.reuters.com/article/us-china-genomics-military-exclusive/exclusive-china-gene-firm-providing-worldwide-covid-tests-worked-with-chinese-military-idUSKBN29Z0HA>.

⁶ *Overview of the Copyright Office*, U.S. Copyright Office (last accessed Feb. 2, 2021), <https://www.copyright.gov/about/>.

⁷ NSCAI staff engagement with Professor Adam Mossoff, Antonin Scalia Law School, George Mason University (Oct. 7, 2020); David J. Kappos, *National Security Consequences of U.S. Patent (In)eligibility*, Morning Consult (Nov. 4, 2019), <https://morningconsult.com/opinions/national-security-consequences-of-u-s-patent-ineligibility/>.

⁸ For example, the Supreme Court's controversial 1980 decision in *Diamond v. Chakrabarty*, which classifies a genetically modified bacterium as a patentable innovation (under Section 101), "was a key factor in spurring the explosive growth in the biotech industry in the ensuing decade in the U.S. The *Chakrabarty* Court's recognition that the products of biotech research are patentable, especially when such products are living organisms or represent the building blocks of life, paved the way for dramatic advances in the life sciences and in medical treatment, such as in cancer research." While the U.S. was the first country to patent genetic modification of living organisms (critical for cancer research), other countries refused to patent such innovations for more than a decade. This led to the U.S. becoming the birthplace of the biotech revolution. Similarly, the Supreme Court's 1981 decision in *Diamond v. Diehr* that an invented process using "a computer program was not automatically an 'abstract idea' or 'algorithm' that precluded patent protection" was key for providing reliable patent rights that enabled the high-tech revolution of the late 20th century. Kevin Madigan & Adam Mossoff, *Turning Gold to Lead: How Patent Eligibility Doctrine Is Undermining U.S. Leadership in Innovation*, George Mason Law Review, Vol. 24 at 942-946 (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2943431.

⁹ Technologies critical to national security interests include AI, microelectronics, 5G telecommunications, quantum computing, and biotechnology. For more information on various U.S. government efforts to define and prioritize critical emerging technologies and the Commission's recommended list of critical emerging technologies, see Chapter 16 of this report and its associated Blueprint for Action. See also *Interim Report and Third Quarter Recommendations*, NSCAI at 138 (Oct. 2020), <https://www.nscai.gov/previous-reports/>. There also is a convergence of technologies with the infusion of AI across all technologies. See Joint Written Testimony of Dr. Eric Schmidt et al. before the House Committee on Armed Services, Subcommittee on Intelligence and Emerging Threats and Capabilities, *Interim Review of the National Security Commission on Artificial Intelligence Effort and Recommendations* (Sept. 17, 2020), <https://docs.house.gov/meetings/AS/AS26/20200917/110996/HHRG-116-AS26-Wstate-SchmidtE-20200917.pdf>.

¹⁰ See *Alice Corp. v. CLS Bank Int'l*, 134 S. Ct. 2347, 2360 (2014) (holding that a computer program for facilitating complex international financial transactions is an abstract idea and cannot be patented); see also *Association for Molecular Pathology v. Myriad Genetics, Inc.*, 133 S. Ct. 2107, 2117 (2013) (holding that isolated DNA for laboratory and medical uses is an unpatentable natural phenomenon); *Mayo Collaborative Services v. Prometheus Laboratories, Inc.*, 566 U.S. 66, 72-73 (2012) (holding that a diagnostic medical treatment for an autoimmune disorder is an unpatentable discovery of a law of nature); *Bilski v. Kappos*, 561 U.S. 593, 609 (2010) (holding that a business method for hedging investment risk is an abstract idea and not a patentable invention); Kevin Madigan & Adam Mossoff, *Turning Gold to Lead: How Patent Eligibility Doctrine Is Undermining U.S. Leadership in Innovation*, *George Mason Law Review*, Vol. 24 at 946-952 (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2943431.

¹¹ A former Chief Judge of the Federal Circuit lamented this uncertainty while testifying before the U.S. Senate Judiciary Committee's Intellectual Property Subcommittee: "It is important for me, as a retired [Federal Circuit] judge, to acknowledge that the courts alone created this problem. ... If I, as a judge with 22 years of experience deciding patent cases on the Federal Circuit's bench, cannot predict outcomes based on case law, how can we expect patent examiners, trial judges, inventors and investors to do so?" See Testimony of Judge Paul R. Michel (Ret.), U.S. Court of Appeals for the Federal Circuit, before the U.S. Senate Committee on the Judiciary, Subcommittee on Intellectual Property, *The State of Patent Eligibility in America: Part I* (June 4, 2019), <https://www.judiciary.senate.gov/imo/media/doc/Michel%20Testimony.pdf>. The U.S. Chamber of Commerce recently observed that uncertainty surrounding patent-eligible subject matter and the viability of biopharmaceutical companies' business models is posing "an existential threat to the United States' position as the undisputed global leader in biopharmaceutical innovation." Art of the Possible: U.S. Chamber International IP Index, U.S. Chamber of Commerce, Global Innovation Policy Center at 10 (2020), https://www.theglobalipcenter.com/wp-content/uploads/2020/02/023881_GIPC_IP_Index_2020_FullReport_A_04b.pdf. The former Director of the USPTO similarly emphasized the importance of certainty to innovation in the U.S.: "[t]o ensure that our nation remains at the forefront of AI and other technologies, we must, among other things, provide a reliable and predictable legal framework to incentivize and protect innovation here at home." See USPTO Responses to Questions for the Record by Senator Tillis, Hon. Andrei Iancu, Under Secretary of Commerce for Intellectual Property and Director of the U.S. Patent and Trademark Office, as Witness, U.S. Senate Committee on the Judiciary, Subcommittee on Intellectual Property, *Oversight of the U.S. Patent and Trademark Office* at 11 (hearing held March 13, 2019, responses submitted Aug. 15, 2019), <https://www.judiciary.senate.gov/imo/media/doc/Iancu%20Responses%20to%20QFRs2.pdf>.

¹² See Joan Farre-Mensa, et al., *What Is a Patent Worth? Evidence from the U.S. Patent "Lottery,"* National Bureau of Economic Research (Dec. 2018), <https://www.nber.org/papers/w23268> (finding an almost double increase in chance of a startup receiving venture capital investments if it has a patent, and further finding this causally linked to a higher rate of success in startups); Stuart J.H. Graham, et al., *High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey*, *Berkeley Technology Law Journal*, Vol. 24, No. 4 at 255-327 (July 4, 2009), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=142904.

¹³ Surveys and industry reports demonstrate that "investment has shifted away from patent-intensive industries." Mark F. Schultz, *The Importance of an Effective and Reliable Patent System to Investment in Critical Technologies*, Alliance for U.S. Startups and Investors for Jobs at 24-37 (July 2020), https://static1.squarespace.com/static/5746149f86db43995675b6bb/t/5f2829980ddf0c536e7132a4/1596467617939/USIJ+Full+Report_Final_2020.pdf. For example, a look at a subset of patent-reliant technologies (core internet networking, wireless communications, internet software, operating system software, semiconductors, pharmaceuticals, drug discovery, surgical devices, and medical supplies) shows a significant decrease in funding, from 21% of total venture capital funding in 2004 to only 3.2% in 2017. *U.S. Startup Company Formation and Venture Capital Funding Trends 2004 to 2017*, Alliance for U.S. Startups and Investors for Jobs at 9 (June 2019), <https://static1.squarespace.com/static/5746149f86db43995675b6bb/t/5d14b7bb46692200012463e0/1561638845187/USIJ+---+U.S.+Startup+Formation+Trends+---+2014-2017.pdf>.

¹⁴ David Taylor, *Patent Eligibility and Investment*, *Cardozo Law Review* at 2055-2056 (2020), <http://cardozolawreview.com/patent-eligibility-and-investment/>.

Blueprint for Action: Chapter 12 - Endnotes

¹⁵ See, e.g., *Hikma Pharmaceuticals USA Inc. v. Vanda Pharmaceuticals Inc.*, No. 18-817 (Jan. 13, 2020) (cert. denied); *Athena Diagnostics, Inc. v. Mayo Collaborative Services, LLC*, No. 19-430 (Jan. 13, 2020) (cert. denied); *HP Inc. v. Berkheimer*, No. 18-415 (Jan. 13, 2020) (cert. denied). In *Athena*, all 12 active judges of the Federal Circuit, the appellate court from which the decision was appealed to the Supreme Court, agreed that the diagnostic methods at issue should be patent eligible, but the majority indicated that they had to find the inventions ineligible for patent protection pursuant to Supreme Court precedent. *Athena Diagnostics, Inc. v. Mayo Collaborative Services, LLC*, No. 19-430 (Jan. 13, 2020) (cert. denied). On Jan. 29, 2021, however, the Supreme Court asked for a response to a petition for certiorari appealing a decision from the Federal Circuit that a drive shaft is not eligible for patent protection because the alleged invention is based on a natural law. *American Axle & Manufacturing Inc. v. Neapco Holdings LLC*, No. 20-891 (Jan. 29, 2021). See also Rebecca Lindhorst, *Two-Stepping Through Alice's Wasteland of Patent-Eligible Subject Matter: Why the Supreme Court Should Replace the Mayo/Alice Test*, Case Western Reserve Law Review at 759 (2019), <https://scholarlycommons.law.case.edu/cgi/viewcontent.cgi?article=4813&context=caselrev>.

¹⁶ In January 2019, the USPTO published the initial framework in a Revised Guidance and requested public comment on the Guidance. See 84 Fed. Reg. 50, *United States Patent and Trademark Office: 2019 Revised Patent Subject Matter Eligibility Guidance*, U.S. Patent and Trademark Office (Jan. 7, 2019), <https://www.federalregister.gov/documents/2019/01/07/2018-28282/2019-revised-patent-subject-matter-eligibility-guidance>. Once the USPTO received comments, it issued an Update to the Guidance: *October 2019 Update: Subject Matter Eligibility*, U.S. Patent and Trademark Office (Oct. 2019), https://www.uspto.gov/sites/default/files/documents/peg_oct_2019_update.pdf. The Revised Guidance and the Update were later incorporated into the newest edition of the USPTO's Manual of Patent Examining Procedure when it was revised in June 2020. See *Manual of Patent Examining Procedure*, United States Patent and Trademark Office at § 2103–2106.07(c) (June 2020), <https://www.uspto.gov/web/offices/pac/mpep/index.html>. Since the USPTO issued the patent eligibility guidance, uncertainty in the examination process has significantly decreased for technologies affected by the *Alice* decision. Office of the Chief Economist, *Adjusting to Alice: USPTO Patent Examination Outcomes After Alice Corp. v. CLS Bank International*, United States Patent and Trademark Office at 6-7 (April 2020), https://www.uspto.gov/sites/default/files/documents/OCE-DH_AdjustingtoAlice.pdf (demonstrating with statistical significance that the Guidance decreased uncertainty as to patent eligibility determinations in the first-action stage of examination by 44% for *Alice*-affected technologies).

¹⁷ Though the USPTO Guidance on patent eligibility applies at the USPTO, the Federal Circuit has held that it is not bound by the Guidance and, if any conflicts arise between it and case precedent from the Federal Circuit and the Supreme Court, precedent will override the Guidance. See *Cleveland Clinic Foundation v. True Health Diagnostics LLC*, 760 F. App'x 1013, 1020 (Fed. Cir. 2019) (non-precedential) (“While we greatly respect the PTO’s expertise on all matters relating to patentability, including patent eligibility, we are not bound by its guidance.”); see also *In re Rudy*, 956 F.3d 1379, 1383 (Fed. Cir. 2020) (precedential) (citing *Cleveland Clinic Foundation*, 760 F. App'x at 1021 (“To the extent the Office Guidance contradicts or does not fully accord with our caselaw, it is our caselaw, and the Supreme Court precedent it is based upon, that must control.”)).

¹⁸ Michael Borella, *The Zombie Apocalypse of Patent Eligibility Reform and a Possible Escape Route*, Patent Docs (Feb. 4, 2020), https://www.patentdocs.org/2020/02/the-zombie-apocalypse-of-patent-eligibility-reform-and-a-possible-escape-route.html?utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+PatentDocs+%28Patent+Docs%29 (citing an interview wherein Senator Thom Tillis, Chairman of the Senate Judiciary Committee’s Subcommittee on Intellectual Property, recognized that his 2019 patent eligibility reform proposal did not have a “path forward” to become a bill in that Congress).

¹⁹ Solely relying on patent counting is not reflective of innovation. See Jonathan Putnam, et al., *Innovative Output in China*, at 32 (Aug. 2020) (pending revision), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3760816.

²⁰ Patrick Thomas & Dewey Murdick, *Patents and Artificial Intelligence: A Primer*, Center for Security and Emerging Technology at 10 (Sept. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Patents-and-Artificial-Intelligence.pdf>.

²¹ Jonathan M. Barnett, *Patent Tigers and Global Innovation*, CATO at 14 (Winter 2019/2020), <https://www.cato.org/sites/cato.org/files/2019-12/v42n4-2.pdf>.

²² WIPO's Patent Cooperation Treaty (PCT) procedure allows inventors to indicate an intent to file patent applications in multiple countries. However, while in the subsequent national phase applications, a third country examines the patent and makes its own determination to grant. Therefore, experts assert that national phase applications are a better indicator for monitoring high-quality patent filings than filings under the PCT system. For information on PCT and national phase process, see *PCT FAQs*, WIPO (April 2020), <https://www.wipo.int/pct/en/faqs/faqs.html>; *WIPO Technology Trends 2019: Artificial Intelligence*, WIPO at 61-63, https://www.wipo.int/edocs/pubdocs/en/wipo_pub_1055.pdf; George Leopold, *China Dominates AI Patent Filings*, *EnterpriseAI* (Aug. 31, 2020), <https://www.enterpriseai.news/2020/08/31/china-dominates-ai-patent-filings/> (“Beijing has become a fierce defender of intellectual property linked to what planners consider a strategic technology”). Although China has a high level of PCT filings, the associated national phase applications are significantly lower. See *Patent Cooperation Treaty Yearly Review 2020: The International Patent System*, WIPO at 50 and 55 (2020), https://www.wipo.int/edocs/pubdocs/en/wipo_pub_901_2020.pdf.

²³ Yuki Okoshi, *China Overtakes U.S. in AI Patent Rankings*, *Nikkei Asia* (March 10, 2019), <https://asia.nikkei.com/Business/Business-trends/China-overtakes-US-in-AI-patent-rankings> (“Chinese companies have surged ahead of their U.S. counterparts on a Nikkei ranking of the top 50 patent filers for artificial intelligence over the past three years, expanding their presence in the world’s most prominent high-tech battleground.”); Andrew Snowdon, *UK Ranked Fourth in the World for Number of Blockchain Patents Filed But Is Falling Behind for AI Patents*, *UHY Hacker Young* (Jan. 21, 2019), <https://www.uhy-uk.com/insights/uk-ranked-fourth-world-number-blockchain-patents-filed-falling-behind-ai-patents> (“New Artificial Intelligence technology developments dominated by Chinese companies”).

²⁴ *5G Technological Leadership*, Hudson Institute at 2 (Dec. 2020), https://s3.amazonaws.com/media.hudson.org/Hudson_5G%20Technological%20Leadership.pdf (“There are important limitations with using patent counting as a measure of innovative output, as economists and statisticians have long recognized. ... This is why economists consider information about the number of patents to be a ‘noisy’ indicator of innovative output. ... What matters is the quality, not the quantity of patents.”); Jonathan Putnam, et al., *Innovative Output in China*, at 32 (Aug. 2020) (pending revision), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3760816; Jonathan M. Barnett, *Patent Tigers and Global Innovation*, *CATO* (Winter 2019/2020), <https://www.cato.org/sites/cato.org/files/2019-12/v42n4-2.pdf>.

²⁵ Michael Mangelson, et al., *Trademarks and Patents in China: The Impact of Non-Market Factors on Filing Trends and IP Systems*, U.S. Patent and Trademark Office at 1 (Jan. 2021), <https://www.uspto.gov/sites/default/files/documents/USPTO-TrademarkPatentsInChina.pdf> (discussing China’s subsidies for trademark and patent application filings); Testimony of Mark Cohen, Senior Counsel on China in the Office of Policy and International Affairs in the United States Patent and Trademark Office, Before House Committee on the Judiciary, *International Antitrust Enforcement: China and Beyond* (June 7, 2016), <https://www.uspto.gov/about-us/news-updates/statement-mark-cohen-house-committee-judiciary> (discussing numerous strategies used by China to increase patent filings).

²⁶ See *Meeting the China Challenge: A New American Strategy for Technology Competition*, Working Group on Science and Technology in U.S.-China Relations at 29 (Nov. 16, 2020), https://asiasociety.org/sites/default/files/inline-files/report_meeting-the-china-challenge_2020.pdf [hereinafter *Meeting the China Challenge*]; Matthew Noble, et al., *Determining Which Companies Are Leading the 5G Race*, *IAM* (July/Aug. 2019), <https://www.twobirds.com/~media/pdfs/news/articles/2019/determining-which-companies-are-leading-the-5g-race.pdf?la=en&hash=8ABA5A7173EEE8FFA612E070C0EA4B4F53CC50DE>. For example, as of February 2020, Huawei and ZTE filed the most number of “standard essential” patents (SEP)s for 5G technologies, but assessments of these filings are critical of the quality of these patents. Jed John Ikoba, *Huawei Has Filed the Most 5G Patents Globally as of February 2020—A Report*, *Gizmochina* (June 2, 2020), <https://www.gizmochina.com/2020/06/02/huawei-has-the-most-5g-standard-essential-patents-globally/>.

Blueprint for Action: Chapter 12 - Endnotes

²⁷ The potential impact of Chinese patent prior art that must be examined at the USPTO can be likened to what is happening to the USPTO trademark application process. An influx of fraudulent trademark applications from China, supported by monetary incentives from the Chinese government, is likely damaging the integrity of the U.S. trademark registration process, including by imposing unpredictability in examination process schedules. *Hearing on Fraudulent Trademarks: How They Undermine the Trademark System and Harm American Consumers and Businesses*, U.S. Senate Committee on the Judiciary, Subcommittee on Intellectual Property (Dec. 3, 2019), <https://www.judiciary.senate.gov/meetings/fraudulent-trademarks-how-they-undermine-the-trademark-system-and-harm-american-consumers-and-businesses>; Barton Beebe & Jeanne C. Fromer, *Are We Running Out of Trademarks? An Empirical Study of Trademark Depletion and Congestion*, Harvard Law Review (Feb. 9, 2018), <https://harvardlawreview.org/2018/02/are-we-running-out-of-trademarks/>; Josh Gerben, *Massive Wave of Fraudulent US Trademark Filings Likely Caused by the Chinese Government Payments*, Gerben (last accessed Jan. 3, 2021), <https://www.gerbenlaw.com/blog/chinese-business-subsidies-linked-to-fraudulent-trademark-filings/>.

²⁸ Jeanne Suchodolski, et al., *Innovation Warfare*, North Carolina Journal of Law & Tech at 201 (Dec. 2020), <https://ncjolt.org/articles/volume-22/volume-22-issue-2/innovation-warfare/>.

²⁹ *Public Views on Artificial Intelligence and Intellectual Property Policy*, U.S. Patent and Trademark Office at iii (Oct. 2020), https://www.uspto.gov/sites/default/files/documents/USPTO_AI-Report_2020-10-07.pdf [hereinafter USPTO AI IP policy report].

³⁰ Richard Vray & Jane Mutimear, *Artificial Intelligence: Navigating the IP Challenges*, Mobile World Live (Feb. 16, 2019), <https://www.mobileworldlive.com/intellectual-property-news/artificial-intelligence-navigating-the-ip-challenges>.

³¹ David Deptula, *The Growing Importance of Data Rights in Defense Acquisition*, Forbes (Oct. 16, 2018), <https://www.forbes.com/sites/davedeptula/2018/10/16/the-growing-importance-of-data-rights-in-defense-acquisition/?sh=165063242a04>.

³² In the USPTO report surveying stakeholders for perspectives on IP policy for AI, “commenters were nearly equally divided between the view that new intellectual property rights were necessary to address AI inventions and the belief that the current U.S. IP framework was adequate to address AI inventions. Generally, however, commenters who did not see the need for new forms of IP rights suggested that developments in AI technology should be monitored to ensure needs were keeping pace with AI technology developments. The majority of opinions requesting new IP rights focused on the need to protect the data associated with AI, particularly ML.” USPTO AI IP policy report at 15; *Id.* at 38 (“[a] smaller number of commenters did suggest a reconsideration of whether additional protections of datasets and databases could be useful to spur investment in high-quality data of vetted/assured provenance.”).

³³ See USPTO AI IP policy report at 15.

³⁴ *Protection of Databases*, European Commission (June 1, 2018), <https://ec.europa.eu/digital-single-market/en/protection-databases>; USPTO AI IP policy report at 38.

³⁵ This includes the Joint Committee on the Research Environment (JCORE).

³⁶ This includes U.S. Customs and Border Protection.

³⁷ This includes the Computer Crime and Intellectual Property Section (CCIPS).

³⁸ Meeting the China Challenge at 16 (“In concert with allies and like-minded countries, the U.S. should investigate, punish, and condemn such acts and identify ways to induce changes in China’s maneuvers through counter-espionage, law enforcement, diplomatic pressure, and professional training in scientific integrity.”).

³⁹ Press Release, The U.S. Department of Commerce, *Statement from Secretary Ross on The Department’s 77 Additions to the Entity List for Human Rights Abuses, Militarization of the South China Sea and U.S. Trade Secret Theft* (Dec. 18, 2020), <https://www.commerce.gov/news/press-releases/2020/12/statement-secretary-ross-departments-77-additions-entity-list-human>.

⁴⁰ Charles Barquist & Maren Laurence, *How a Biden Administration Would Shape IP Policy*, Law 360 (Oct. 19, 2020), <https://www.law360.com/articles/1319716/how-a-biden-administration-would-shape-ip-policy>; Sean Lyngaas, *As China Tensions Mount, U.S. Officials Outline Efforts to Combat Economic Espionage*, CyberScoop (Dec. 12, 2018), <https://www.cyberscoop.com/china-tensions-mount-u-s-officials-outline-efforts-combat-economic-espionage/>; see also 18 U.S.C. § 1831 (regarding economic espionage); 18 U.S.C. §1832 (regarding theft of trade secrets).

⁴¹ Robert Bahr, *Decision on Petition: Application No. 16/524,350*, U.S. Patent and Trademark Office (2020), https://www.uspto.gov/sites/default/files/documents/16524350_22apr2020.pdf.

⁴² Consistent with U.S. policy that an inventor must be a human “natural person,” in January 2020 the European Patent Office (EPO) and the UK Intellectual Property Office (UKIPO) rejected two patent applications that identified the AI machine as the inventor. The EPO and UKIPO found that the applications met the requirements for patentability, but they rejected the applications because the inventor was not a “human being.” See Emma Woollacott, *European Patent Office Rejects World’s First AI Inventor*, Forbes (Jan. 3, 2020), <https://www.forbes.com/sites/emmawoollacott/2020/01/03/european-patent-office-rejects-worlds-first-ai-inventor/?sh=2915e17d5cd0>; Angela Chen, *Can an AI Be an Inventor? Not Yet*, MIT Technology Review (Jan. 8, 2020), <https://www.technologyreview.com/2020/01/08/102298/ai-inventor-patent-dabus-intellectual-property-uk-european-patent-office-law/>; *EPO Provides Reasoning for Rejecting Patent Applications Citing AI as Inventor*, IPWatchdog (Jan. 28, 2020), <https://www.ipwatchdog.com/2020/01/28/epo-provides-reasoning-rejecting-patent-applications-citing-ai-inventor/id=118280/>.

⁴³ USPTO AI IP policy report at ii-iii.

⁴⁴ See the Chapter 15 Blueprint for Action and its associated Annex for more details on the proposed critical areas for international alignment for the Emerging Technology Coalition. Critical Area No. 4, as detailed in the Blueprint for Action and Annex, is “Promoting and protecting innovation, including through intellectual property alignment.” Recognizing the importance of IP to promote and protect innovation, the critical area proposes coordination on assistance to nations in developing strong and aligned IP regimes, coordinated efforts to stop IP theft and counter-cyberespionage, and aligning on a mutual agenda within IP-related multilateral forums.

⁴⁵ “To maintain our technological leadership, the United States must seek to broaden our intellectual property ecosystem demographically, geographically, and economically.” *Expanding Innovation*, USPTO (last accessed Jan. 3, 2021), <https://www.uspto.gov/initiatives/expanding-innovation> (quoting USPTO Director Andrei Iancu).

⁴⁶ *Remarks by Commerce Secretary Wilbur L. Ross at the First Meeting of the National Council for Expanding American Innovation*, U.S. Department of Commerce (Sept. 14, 2020), <https://www.commerce.gov/news/speeches/2020/09/remarks-commerce-secretary-wilbur-l-ross-first-meeting-national-council>; *Support the National Council for Expanding American Innovation*, USPTO (last accessed Jan. 3, 2021), <https://www.uspto.gov/initiatives/expanding-innovation/national-council-expanding-innovation/support-national-council>.

⁴⁷ “A significant proportion of lawyers are advising clients with products in the global market to patent in China, Germany, and even the U.K. instead of the U.S. The U.S. is losing the fight to be the major center of patents, investment, and tech because it is easier and less expensive for companies to file and ensure their patents are enforced in other countries than in the U.S.” NSCAI staff engagement with Robert Taylor, owner of RPT Legal Strategies, PC (Oct. 8, 2020).

⁴⁸ Through the standards-setting process, standards-setting bodies (e.g., ISO, IEC, IEEE, ITU, and others) often require that patent owners self-identify patents that *may be* deemed essential in a future standard. This requirement aims to ensure transparency and often requires commitments by these patent owners to license their patents fairly, reasonably, and non-discriminatorily. However, these standards-setting bodies do not assess whether a patent is essential or not, leaving these determinations to private companies negotiating licenses or, if there is a dispute, to courts. See *IEEE SA Standards Board Bylaws*, IEEE, <https://standards.ieee.org/about/policies/bylaws/sect6-7.html#loa>.

Blueprint for Action: Chapter 12 - Endnotes

⁴⁹ See Chapter 15 of this report and its associated Blueprint for Action for the coordinated U.S. national plan to support international technology efforts and its first component on shaping international technical standards. Also see the Chapter 15 Annex for more details on proposed international technical standards-setting recommendations for NIST, the Department of State, and other critical Departments and Agencies. NSCAI recommends that the U.S. government provide greater attention to and resourcing for international technical standardization efforts; increase interagency coordination on AI-related standards-setting; strengthen partnerships and collaboration with the private sector, particularly through a federal advisory committee and a grant program to enable small and medium-sized U.S. AI companies to participate in international standardization efforts; and increase international alignment with key partners and allies. See also Meeting the China Challenge at 27.

⁵⁰ Dai Hong, the director of China's National Standardization Committee's Industrial Standards Department, stated in January 2018, as the research for China Standards 2035 was launched: "In today's world, industry, technology, and innovation are developing rapidly. The new generation of information technology industry represented by artificial intelligence, big data, cloud computing, etc. is emergent. International technology research and development and patent distribution have not yet been completed. Global technical standards are still being formed. This offers the opportunity to realize the transcendence of China's industry and standards." See translated quote from January 20, 2018, on the China News Network in Emily de la Bruyère & Nathan Picarsic, *China Standards 2035: Beijing's Platform Geopolitics and 'Standardization Work in 2020,'* Horizon Advisory at 6 (April 2020), <https://www.horizonadvisory.org/china-standards-2035-first-report>. Additionally, the Guangdong High People's Court published an October 2013 opinion piece that argued "for Chinese enterprises to make a revival, there is only one road to take: strengthen our capacity for innovation, and only by gaining control over SEPs can Chinese companies avoid being 'led by the nose.'" It cited Chief Judge Qiu Yongqing, who ruled against the U.S. firm InterDigital in its lawsuit against Huawei and argued that "Chinese enterprises should bravely employ anti-monopoly lawsuits to break technology barriers and win space for development." See David Cohen & Douglas Clark, *China's Anti-Monopoly Law as a Weapon Against Foreigners*, IAM-media (Nov./Dec. 2018), https://kidonip.com/wp-content/uploads/2018/11/IAM92_China-anti-monopoly_section_0.pdf.

⁵¹ Jeanne Suchodolski, et al., *Innovation Warfare*, North Carolina Journal of Law & Tech at 201 n. 130 (Dec. 2020), <https://ncjolt.org/articles/volume-22/volume-22-issue-2/innovation-warfare/> (China's firms recognize the strategic importance of standards-setting activities and that participation in those forums provides the legal means to both access and influence developing technologies). "In recent years the PRC government decided that promoting Chinese standards in global standards bodies via the work of Huawei and other Chinese companies is key to realizing techno-nationalist goals for technological ascension. Viewed in this context, Huawei is in the vanguard of the Chinese effort to establish dominance in both the number and significance of Chinese patents that are deemed "standard essential" to 5G standards ... it is in the U.S. interest to deflect Beijing's attempt to dominate the standard-setting process." See Meeting the China Challenge at 29. See also Matthew Noble, et al., *Determining Which Companies Are Leading the 5G Race*, IAM (July/August 2019), <https://www.twobirds.com/~media/pdfs/news/articles/2019/determining-which-companies-are-leading-the-5g-race.pdf>.

⁵² Over-declaration is already present in 5G. See Matthew Noble, et al., *Determining Which Companies Are Leading the 5G Race*, IAM (July/August 2019), <https://www.twobirds.com/~media/pdfs/news/articles/2019/determining-which-companies-are-leading-the-5g-race.pdf>.

⁵³ Countries are increasingly seeking to attract inventors by setting favorable global royalty rates (see the U.K.'s decision in *Unwired Planet v. Huawei*) or by controlling the jurisdiction in which companies may file for injunctive relief or pursue litigation. For example, licensing disputes have recently led to additional satellite litigation involving broader issues of international law and comity between China and other legal jurisdictions. Experts predict disputes to increase and warn of cycle of anti-suit, "anti-antisuit," and "anti-anti-antisuit" injunctions. See Mark Cohen, *Wuhan and Anti-Suit Injunction*, China IPR Blog (Dec. 28, 2020), <https://chinaipr.com/2020/12/28/wuhan-and-anti-suit-injunctions/>; Dani Kass, *FRAND Rate 'Nightmare' Raises Call For International Tribunal*, Law360 (Jan. 14, 2021), <https://www.law360.com/articles/1343824/frand-rate-nightmare-raises-call-for-international-tribunal/>; Michael Renaud, et al., *Key Considerations for Global SEP Litigation—Part 1*, National Law Review (Oct. 30, 2019), <https://www.natlawreview.com/article/key-considerations-global-sep-litigation-part-1>; Michael Renaud, et al., *Key Considerations for Global SEP Litigation—Part 2*, National Law Review (Nov. 5, 2019), <https://www.natlawreview.com/article/key-considerations-global-sep-litigation-part-2>; Zhao Qishan & Lu Zhe, *Statistics of Chinese SEP Cases in 2011-2019*, LexField (2020), <https://chinaipr2.files.wordpress.com/2020/07/statistics-of-chinese-sep-cases-in-2011-2019-lexfield9892.pdf>.

Chapter 13:

Microelectronics

Blueprint for Action



Regaining microelectronics leadership requires meeting an explicit objective: Stay at least two generations ahead of China in state-of-the-art microelectronics and maintain multiple sources of cutting-edge microelectronics fabrication in the United States. To do this, the Executive Branch must prepare and implement a National Microelectronics Strategy while Congress simultaneously institutes new tax credits, subsidizes the construction of semiconductor manufacturing facilities, and grows federal microelectronics R&D and infrastructure funding. Achieving this goal will require roughly \$30 billion in additional federal funding, but these funds should attract more than five times as much private-sector investment. Additional federal funding on this scale will likely boost economic activity domestically and could add more than \$100 billion to U.S. gross domestic product (GDP).¹ Inside the U.S. government, agencies must also expand access to trustworthy, high-performance microelectronic components by shifting from serial to concurrent development of hardware and software to catch up to the commercial sector and make use of new microelectronics produced in the United States.

Five-Year Microelectronics Funding

Category	Amount
Federal Grants for Microelectronics Manufacturing*	\$3 billion per project (\$15 billion total)
Microelectronics R&D	\$12 billion
Microelectronics Infrastructure	\$7 billion
DoD Trusted & Assured Microelectronics	\$0.5 billion
Total	\$35 billion

Recommendation: Issue an Executive Order on Microelectronics Strategy and Leadership

Recommendation

The United States needs a National Microelectronics Strategy to coordinate semiconductor policy, funding, and incentives within the Executive Branch and externally with industry and academia.

Actions for the President:

- **Issue an Executive Order on Microelectronics National Strategy and Leadership.**
 - o The first step in rebuilding microelectronics leadership is clearly stating that it is a Presidential priority to stay at least two generations ahead and maintain multiple sources of cutting-edge microelectronics fabrication in the United States. The Administration should also highlight the importance of the legislatively required National Microelectronics Strategy and create a durable structure for its development, implementation, and revision by issuing an Executive Order requiring the National Defense Authorization Act (NDAA)-mandated Subcommittee on Microelectronics Leadership to lead a process to develop a clear federal strategy for microelectronics leadership. Draft text to inform the development of an Executive Order for this purpose is included as an Annex to this Blueprint for Action.

Recommendation: Revitalize Domestic Microelectronics Fabrication

Recommendation

Existing U.S. incentives offset the cost of semiconductor foundry construction attributable to capital expenses, operating expenses, and taxes by 10% to 15%.² Yet additional tax credits and subsidies are needed to make the United States a globally competitive market for semiconductor manufacturing, especially leading-edge logic facilities. Other leading semiconductor manufacturing nations such as South Korea, Taiwan, and Singapore offer 25% to 30% cost reduction, roughly double what the United States currently offers.³ This gap in incentives is one driving factor behind the lack of an advanced logic merchant foundry in the United States. Closing the gap will encourage U.S. firms to construct facilities domestically while also attracting foreign firms. In fact, a program of the size described here is projected to attract roughly 14 new fabs in the United States over 10 years.⁴ Additionally, increasing demand in the United States for high-end semiconductor manufacturing equipment (SME) will create new business opportunities for SME manufacturers from allied countries, particularly Japan and the Netherlands, which could increase their governments' willingness to align their export control policies with U.S. policies prohibiting the export of such equipment to China.⁵ A refundable investment tax credit should be instituted in combination with funding for federal grants for the expansion, construction, and modernization of SME authorized in the NDAA.⁶

Action for Congress:

- **Create a 40% refundable investment tax credit for domestic semiconductor manufacturing.**

- o Congress should pass legislation establishing a 40% refundable federal investment tax credit for semiconductor manufacturing facilities and equipment required to produce state-of-the-art logic chips. This incentive would reduce a semiconductor firm's tax bill by 40% on SME and facilities through 2024, followed by reduced tax credit rates of 30% and 20%, respectively, through 2025 and 2026. Although introduced as part of the Creating Helpful Incentives to Produce Semiconductors for America Act (CHIPS for America Act), Congress has not yet passed legislation establishing this credit.⁷
- **Appropriate funding authorized in the FY 2021 NDAA for domestic semiconductor manufacturing incentives, including matching funds for semiconductor fabrication facilities.**
 - o The FY 2021 NDAA authorizes the Secretary of Commerce to establish a Federal financial assistance program to incentivize investment in facilities and equipment in the United States for semiconductor fabrication, assembly, testing, advanced packaging, or R&D. Under the program, the Secretary may authorize up to \$3 billion per project to finance the construction, expansion, or modernization of facilities and equipment for semiconductor manufacturing. Larger subsidies are also permitted if the project significantly increases the proportion of semiconductors relevant for national security and economic competitiveness that can be met through reliable domestic production. However, this judgment requires the concurrence of the Secretary of Defense and Director of National Intelligence.
 - o Although authorized in the FY 2021 NDAA, funds have not yet been appropriated toward this program. Congress should appropriate at least \$15 billion to subsidize several facilities in the United States to meet the end goal of multiple state-of-the-art sources for domestic fabrication.

Recommendation

Recommendation: Double Down on Funding for Research and Infrastructure to Lead the Next Generation of Microelectronics

Four research arms of the U.S. government focused on medium- and long-term microelectronics breakthroughs through engagement with academia and industry are the Department of Energy (DOE), the Defense Advanced Research Projects Agency (DARPA), the National Science Foundation (NSF), and the Department of Commerce. Their suite of existing programs, such as DARPA's Electronics Resurgence Initiative (ERI), are targeting the right research areas but must be expanded by an order of magnitude to achieve the necessary breakthroughs to maintain U.S. competitiveness. Additional funding should support not only research projects, but also the capital-intensive infrastructure for microelectronics development, including the National Semiconductor Technology Center and advanced packaging prototyping programs authorized in the FY 2021 NDAA. In line with the existing focus areas of these programs and the Commission's prior recommendations, funding should pursue breakthroughs in promising technologies such as 3D chip stacking, photonics, carbon nanotubes, gallium nitride transistors, domain-specific hardware architectures, electronic design automation, and cryogenic computing. In particular, funding should prioritize the development of manufacturing equipment and tools to reach 3nm and beyond at production scale. However, this funding should not solely be directed to classical computing technologies. The U.S. government should also support efforts to research and develop hybrid quantum-classical techniques that

leverage noisy intermediate-scale quantum computers. The Commission offers detailed recommendations on this subject in Chapter 16 of this report.

Action for Congress:

- **Appropriate \$1.1 billion for semiconductor R&D in FY 2022 and continue increasing funding over the next five years for a total of \$12 billion**
 - Congress should appropriate an additional \$1.1 billion in FY 2022. Consistent with the amounts in the CHIPS for America Act, this funding should include \$400 million for DARPA ERI, \$300 million for NSF semiconductor research, and \$400 million for DOE semiconductor research. These funding levels should be grown over the following five years to roughly \$7 billion per year and \$12 billion total. Recognizing it will take time to build capacity among agencies to administer programs at the necessary scale, these amounts should start at funding levels that can be absorbed by agencies and ramped up over time.
- **Appropriate \$1 billion in FY 2022 and \$5 billion total over five years for the Advanced Packaging National Manufacturing Program.**
 - Novel packaging techniques such as heterogeneous integration and 3D stacking—combined with domain-specific architectures—will be critical to the development of artificial intelligence (AI) as traditional architectures of silicon-based chipsets encounter diminishing marginal performance improvements. Congress should also appropriate \$1 billion in initial FY 2022 funding to establish the Advanced Packaging National Manufacturing Program led by the National Institute of Standards and Technology (NIST), as authorized by the FY 2021 NDAA.⁸ This funding should be continued through FY 2027 for a total of \$5 billion.
- **Appropriate \$100 million in FY 2022 and \$2 billion over five years to establish the National Semiconductor Technology Center.**
 - A National Semiconductor Technology Center would serve as a microelectronics research hub while also conducting prototyping of advanced semiconductors in partnership with the private sector. Early-stage semiconductor startups currently face difficulties scaling due to the high costs of microelectronics design and fabrication. The incubator component of the center could provide resources to promising, early-stage microelectronics startups while also giving them access to fabrication facilities, design tools, and shared intellectual property (IP) to assist with early-stage development costs. It could also partner with the U.S. International Development Finance Corporation (DFC) to provide loan guarantees to microelectronics firms if DFC's authorities are expanded and extended to rebuild domestic supply chains for a broader range of strategic emerging technologies.⁹ This laboratory could grow into a center of expertise in high-performing, trusted microelectronics, ensuring continued U.S. leadership in this field over the ensuing years.

Recommendation: Continue DoD's Trusted Microelectronics Program and Adopt Agile Hardware Development

Recommendation

Semiconductor manufacturing has moved offshore, expanding threat vectors to hardware security and leaving the U.S. government unable to trust sensitive electronic components it needs for defense systems. And while the U.S. government is now recognizing that it

must take steps to adopt modern software practices, there has been less attention on incorporating hardware into the agile development process. Both issues require attention from the Department of Defense (DoD) and other government agencies. The U.S. government needs to inject security and agility into its microelectronics acquisition and development process to leverage the best technology possible for defense systems.

Actions for the Department of Defense:

- **Continue growing the Trusted & Assured Microelectronics Program to include AI-enabling hardware.**

- o DoD's Trusted and Assured Microelectronics research, development, test, and evaluation (RDT&E) funding has grown to more than \$500 million annually for advanced component development and prototyping and system development and demonstration.¹⁰ These programs improve access to advanced packaging and testing; support the development of quantifiable assurance and secure design; develop foundry access standards; expand access to non-complementary metal oxide semiconductor state-of-the-art microelectronics; support disruptive R&D; and promote education and workforce development. These are foundational microelectronics capabilities that will also enable the development and application of AI and machine learning (ML) capabilities across national security mission areas. In FY 2021 and beyond, USD(R&E) should expand the program to focus on developing AI-enabling capabilities and apply \$50 million of funding toward developing AI multi-chip packages.

- **Shift to a more agile approach to hardware development and procurement.**

- o Just as agile development has transformed software, there is an opportunity to bring agile hardware design practices to speed development cycles, lower costs, and increase performance. Rather than designing through a serial process, the commercial sector has developed best practices to integrate hardware and software development processes concurrently. While DoD has made strides in agile software development, it remains behind the commercial sector in applying these lessons to hardware. Broader adoption of hardware emulation and moving to a common and secure design environment for the chip, package, and board would also accelerate system development and improve security. This requires the combined efforts of USD(R&E) and USD(A&S) to continue improving software acquisition and development practices to incorporate hardware.

Blueprint for Action: Chapter 13 - Endnotes

¹ *Sparking Innovation: How Federal Investment in Semiconductor R&D Spurs U.S. Economic Growth and Job Creation*, Semiconductor Industry Association at 2 (June 2020), https://www.semiconductors.org/wp-content/uploads/2020/06/SIA_Sparking-Innovation2020.pdf; *Semiconductor Incentives*, Semiconductor Industry Association at 2 (Oct. 9, 2020), <https://www.semiconductors.org/wp-content/uploads/2020/10/Incentives-Infographic-2020.pdf>.

² Antonio Varas, et al., *Government Incentives and US Competitiveness in Semiconductor Manufacturing*, Boston Consulting Group and Semiconductor Industry Association at 19 (Sept. 2020), <https://web-assets.bcg.com/27/cf/9fa28eeb43649ef8674fe764726d/bcg-government-incentives-and-us-competitiveness-in-semiconductor-manufacturing-sep-2020.pdf>.

³ *Id.*

⁴ *Id.*

⁵ See Chapter 14 of this report for additional details regarding export controls on SME.

⁶ Total matching funding will vary based on the number of projects approved but should have a ceiling of at least \$10 billion to \$15 billion.

⁷ S. 3933, 116th Cong. (2020); H.R. 7178, 116th Cong. (2020).

⁸ Pub. L. 116-283, sec. 9906, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁹ See the Chapter 16 Blueprint for Action for further details on extending and expanding DFC's loan guarantee program through executive action.

¹⁰ Pub. L. 116-260, Division C, Department of Defense Appropriations Act (2021), <https://docs.house.gov/billsthisweek/20201221/BILLS-116RCP68-JES-DIVISION-C.pdf>.

Chapter 13 Annex: Executive Order on Microelectronics Strategy

By the authority vested in me as President by the Constitution and the laws of the United States of America, including section 9906 of the National Defense Authorization Act (NDAA) for Fiscal Year 2021 (Public Law 116-283), it is hereby ordered as follows:

Section 1. Findings. The United States relies heavily on imports of certain microelectronics that are vital to the Nation's security and economic prosperity. This dependency on semiconductor imports creates strategic economic and military vulnerabilities to supply chain disruptions for electronics, including adverse foreign government actions and natural disasters. Despite tremendous expertise in microelectronics research, development, and innovation across the country, the United States is limited by a lack of domestically located semiconductor fabrication facilities, especially for state-of-the-art semiconductors. This limitation compounds the risk that the United States may be outpaced in microelectronics design and fabrication. Focusing the efforts of the United States Government, industry, and academia to develop domestic microelectronics fabrication facilities will reduce the Nation's dependence on imports, preserve U.S. leadership in technological innovation, support job creation, strengthen national security and balance of trade, and enhance the technological superiority and readiness of the Armed Forces, which are important consumers of advanced microelectronics.

Section 2. Policy. To maintain the Nation's security and economic prosperity, it shall be the policy of the United States to stay at least two generations ahead of potential adversaries in state-of-the-art microelectronics and maintain multiple sources of cutting-edge microelectronics fabrication in the United States.

Section 3. Establishment of Subcommittee on Microelectronics Leadership.

(a) There is hereby established in the National Science and Technology Council a subcommittee on matters relating to leadership and competitiveness of the United States in microelectronics technology and innovation to be named the Subcommittee on Microelectronics Leadership (Subcommittee).

(b) The Subcommittee shall be composed of the following members:

- (i) The Secretary of Commerce, who shall be Chair of the Subcommittee;
- (ii) The Secretary of State;
- (iii) The Secretary of Defense;
- (iv) The Secretary of Energy;
- (v) The Secretary of Homeland Security;

- (vi) The Director of the Office of Management and Budget;
- (vii) The United States Trade Representative;
- (viii) The Director of National Intelligence;
- (ix) The Director of the National Science Foundation;
- (x) The Assistant to the President for Science and Technology;
- (xi) The Assistant to the President for Technology Competitiveness;
- (xii) The Assistant to the President for National Security Affairs;
- (xiii) The Assistant to the President for Economic Policy;
- (xiv) The Assistant to the President for Domestic Policy; and

(xv) The heads of other executive departments and agencies and other senior officials within the Executive Office of the President, as determined by the Chair.

(c) Sunset. The Subcommittee shall terminate on January 1, 2031.

Section 4. Functions of the Subcommittee on Microelectronics Leadership.

Consistent with applicable law, the Subcommittee shall:

- (a) advise the President on matters involving policy affecting microelectronics;
- (b) develop, within 270 days of the date of this order, and no less than once every five years thereafter, a National Strategy on Microelectronics Research, Development, Manufacturing, and Supply Chain Security (Strategy), which shall address the following elements:
 - (i) methods to accelerate the domestic development and production of microelectronics and strengthen the domestic microelectronics workforce;
 - (ii) methods to ensure that the United States is a global leader in the field of microelectronics research and development;
 - (iii) activities that may be carried out to strengthen engagement and outreach between Federal agencies and industry, academia, and international partners of the United States on issues relating to microelectronics;

(iv) priorities for research and development to accelerate the advancement and adoption of innovative microelectronics and new uses of microelectronics and components;

(v) the role of diplomacy and trade in maintaining the position of the United States as a global leader in the field of microelectronics;

(vi) the potential role of a Federal laboratory, center, or incubator exclusively focused on the research and development of microelectronics, as described in section 231(b)(15) of the NDAA for Fiscal Year 2017 (as added by section 276 of the NDAA for Fiscal Year 2021) in carrying out the Strategy; and

(vii) such other activities as the Subcommittee determines may be appropriate to overcome future challenges to the innovation, competitiveness, and supply chain integrity of the United States in the field of microelectronics; and

(c) coordinate the policymaking process with respect to microelectronics-related research, development, manufacturing, and supply chain security activities and budgets of Federal agencies and ensure such activities are consistent with the Strategy required by this section.

Section 5. General Provisions. (a) If any provision of this order or the application of such provision is held to be invalid, the remainder of this order and other dissimilar applications of such provision shall not be affected.

(b) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.

(c) Nothing in this order shall be construed to impair or otherwise affect:

(i) the authority granted by law to an executive department or agency, or the head thereof; or

(ii) the functions of the Director of the Office of Management and Budget relating to budgetary, administrative, or legislative proposals.

(d) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.

Chapter 14:

Technology Protection

Blueprint for Action

This Blueprint for Action provides detail for how the United States must craft technology protection policies to ensure it retains existing advantages in technology areas with national security applicability but avoids stifling innovation. U.S. research, entrepreneurship, and talent development remain the key ingredients of success. However, as dual-use technologies become more important to U.S. national security, the margin of U.S. technological advantage narrows, and foreign efforts to acquire American know-how and technology increase, the United States must also reexamine how it can protect its commercial and academic ecosystem from foreign exploitation. The United States faces substantial challenges in adapting its technology protection regime to address threats related to emerging, dual-use technologies such as artificial intelligence (AI) without hindering the free flow of commerce or its open research environment, both of which are systemic U.S. strengths. This Blueprint for Action proposes reforms for (1) modernizing export controls and investment screening and (2) protecting the U.S. research environment in ways which are consistent with U.S. national security, commercial interests, and values.

Modernizing Export Controls and Investment Screening

How the U.S. Government regulates competitors' access to sophisticated U.S. technologies with national security applications will be one of the principal challenges of current and future geoeconomic competition. The United States must modernize its export control and investment screening regimes to better address the challenges posed by dual-use emerging technologies, to include AI. These reforms are necessary to allow the government to implement technology protection policies in ways which maximize their impact on the military capabilities of U.S. strategic competitors and minimize any resulting harms to U.S. industry.

Recommendation

Recommendation: Clearly State the Overarching Principles to Guide Future U.S. Dual-Use Technology Protection Policies

The U.S. Government must clearly state the principles that will guide future U.S. decisions regarding policies to protect critical technologies. This will enable more consistent and cohesive technology protection policies and provide clarity to industry regarding how the government intends to utilize these regulatory tools in the current competitive environment, thereby reducing uncertainty for U.S. businesses. No such framework currently exists.

Action for the President:

- **Issue an Executive Order outlining the principles which will guide U.S. policies for protecting dual-use technologies.¹**
 - o The President should issue an Executive Order to clarify guiding principles which will guide U.S. policies to protect critical dual-use technologies, including AI. The Executive Order should include the following guiding principles:
 - U.S. technology controls will not supplant investment and innovation.
 - U.S. strategies to promote and protect U.S. technology leadership will be integrated and mutually reinforcing.
 - The United States will be judicious in applying export controls to AI-related technologies, targeting discrete chokepoints and coordinating policies with allies.
 - The United States will broaden investment screening to protect AI-related technologies.

Recommendation: Enhance U.S. Capacity to Carry Out Effective Technology Protection Policies

Recommendation

Departments and agencies responsible for protecting U.S. technologies lack the organizational and technical capacity to design and implement effective policies to prevent the transfer of the national security–sensitive components of emerging technologies such as AI. They suffer from a dearth of technical talent needed to identify effective new policies and lack the analytical capacity to enforce their policies efficiently, especially on dual-use goods. Filling these gaps in key elements of the Executive Branch—particularly in the Departments of Commerce, the Treasury, and State—will enhance the government’s ability to craft targeted export controls that have the greatest strategic impact and pose the least harm to U.S. competitiveness.

Actions for the Department of Commerce:

- **Designate a network of FFRDCs and UARCs to serve as a shared technical resource on export controls.²**
 - o To deepen its internal technical expertise, the Department of Commerce should establish a network within existing federally funded research and development centers (FFRDCs) and university–affiliated research centers (UARCs) to provide technical expertise to all departments and agencies for issues relating to export controls on emerging technologies. This network should be coordinated by the Department of Commerce and encompass a regional distribution of FFRDCs and UARCs that are either located in U.S. technology hubs or have significant expertise in emerging technologies.
 - o As an initial step, the Department of Commerce should identify the FFRDCs and UARCs with existing expertise in emerging technologies under consideration for export controls. This should be followed by a request for funding in the Fiscal Year (FY) 2022 President’s Budget to support and expand work of FFRDCs and UARCs focusing on export controls.

- **Require all new technology protection rules on emerging technologies to be coordinated with existing technical advisory groups that include outside experts.**³

- o The Secretary of Commerce should require that the Bureau of Industry and Security (BIS) solicit and receive feedback on any proposed controls on emerging or foundational technologies, to include proposed rules and regulations, from the Emerging Technology Technical Advisory Committee (ETTAC) and any other relevant technical advisory groups.⁴ More frequent and effective use of such existing advisory committees would provide flexible technical expertise to key departments, help prevent publishing counterproductive controls, and ensure that policymakers hear the perspective of industry and academia before controls go into effect.

Actions for the Departments of Commerce, the Treasury, and State:

- **Expedite and automate export licensing and CFIUS filing processes.**⁵

- o The Departments of Commerce and the Treasury should partner with FFRDCs, UARCs, and other contracted entities to build an integrated, smart system for analyzing export license applications and filings with the Committee on Foreign Investment in the United States (CFIUS). This system should utilize AI to conduct a preliminary analysis of filings and attempt to score levels of risk before human review. In the near term, this would help identify which transactions are very low risk and which are very high risk to aid subsequent human review. In the longer term, it could prove more accurate than human review and make decisions without human involvement, allowing for precise, rapid, and less labor-intensive reviews.

- **Encourage allies to implement legal reforms authorizing them to implement unilateral export controls and enhance investment screening procedures.**

- o The Departments of State and Commerce must urge all allies which have not already done so to pass domestic legislation to overhaul their export control regimes, increasing their bureaucratic capacity and providing them the authorities to implement unilateral export controls. Currently, many allies lack such domestic legal authorities and instead defer all decisions about regulations to multilateral organizations such as the Wassenaar Arrangement and the European Union.⁶ These reforms are needed to allow allies to implement targeted, rapid, and effective export controls on emerging dual-use technologies, which are evolving quickly. Technology protection regimes on globally available products are only as strong as their weakest link, necessitating U.S. cooperation with allies and strong allied regulatory capacity. This builds on existing work, which has been productive and should continue with an immediate focus on countries that have a strong domestic emerging technology base and weak regulatory regimes.⁷
- o The Departments of State and the Treasury should expedite efforts to enhance the investment screening capabilities of close allies and partners. Existing efforts have shown some success but now require increased urgency, given the threats allies face from adversarial capital and the U.S. desire to exempt some firms in allied nations from certain CFIUS requirements.⁸ State and the Treasury should also regularly share data about patterns in investment flows in the United States and allied countries to assist allied efforts to block predatory investments and illustrate the nature of the threat.

- **Ensure that the offices responsible for export controls and investment screening policies have sufficient resources and technical capacity.**

- o The Departments of Commerce, the Treasury, and State must ensure that the offices responsible for designing and implementing export controls and investment screening provisions on emerging technologies are sufficiently resourced and have sufficient technical capacity. Agencies should rely on external sources such as FFRDCs, UARCs, and advisory boards for deep technical expertise on particular technologies. However, they also must ensure that the offices principally responsible for managing the policy processes regarding controls on these technologies have adequate staffing, resources, and baseline technical capacity to keep pace with the rapidly evolving security challenges associated with dual-use technologies.

Recommendation: Identify “Emerging” and “Foundational” Technologies Which Must Be Controlled, as Required by the Export Control Reform Act of 2018

Recommendation

The Export Control Reform Act of 2018 (ECRA) and the Foreign Investment Risk Review Modernization Act of 2018 (FIRRMA) are intended to overhaul the U.S. export control and investment screening regimes to better accommodate emerging technologies. ECRA requires the Department of Commerce to develop a regular, formal interagency process to identify “emerging and foundational technologies that ... are essential to the national security of the United States,” and are not otherwise controlled.⁹ Any such technologies identified by Commerce become subject to U.S. export controls, and any foreign investment in a U.S. company which “produces, designs, tests, manufactures, fabricates, or develops” one or more such technologies must be reviewed by CFIUS.¹⁰ This list must be distinct from efforts within the Commission-proposed National Technology Strategy (NTS) to define emerging technologies key to U.S. national competitiveness and national security. The ECRA list must be more narrowly defined and focused only on specific technologies for which export controls are necessary, whereas the TCC and NTS’ focus should be on identifying broader technologies and particular platforms in which continued U.S. leadership is essential.

However, as of March 2021, the Department of Commerce has yet to identify a single emerging or foundational technology as mandated by ECRA. While there is reason to be judicious in developing this list, given its implications on U.S. industry, and Commerce faces legitimate capacity and resourcing limitations, the magnitude of the delay is unacceptable. The delay has garnered bipartisan criticism, created uncertainty for firms working in fields that could be labeled as emerging or foundational technologies, and delayed the government’s ability to either control the export of, or more importantly gain insight into transactions involving, critical technologies that are not otherwise controlled.¹¹

Identifying this list of technologies is critical to enabling the United States to fully implement both ECRA and FIRRMA. As ECRA and FIRRMA are structured, until the Department of Commerce defines a technology which is not otherwise controlled as “emerging and foundational” as part of this review process, with rare exceptions CFIUS cannot require foreign companies to disclose non-controlling investments in U.S. technology

firms. Although the Commission also recommends breaking CFIUS' reliance on this ECRA list for mandatory disclosures (see recommendations on reforming CFIUS for emerging technology competition, below), currently Commerce's delay in identifying such technologies is hindering the full implementation of both ECRA and FIRRMA.

Action for the Department of Commerce:

- **Direct the Bureau of Industry and Security to develop proposed rules containing initial lists of both “emerging” and “foundational” technologies by December 31, 2021.¹²**

- o The Secretary of Commerce should direct the BIS to work with the U.S. interagency to develop initial versions of the lists of “emerging” and “foundational” technologies by December 31, 2021. Beyond 2021, these lists should be regularly revised in an iterative manner to meet ECRA's mandate to Commerce to continually refine the lists. As part of this iterative review process, Commerce must also regularly engage with industry as technologies develop and mature. Finalizing initial versions of these lists, if properly scoped and defined, would control critical technologies, clarify to industry how Commerce intends to implement ECRA, and ensure that such technologies are included within CFIUS.

Recommendation

Recommendation: Reform CFIUS for Emerging Technology Competition

CFIUS is not currently postured to address the range of threats that the United States faces from adversarial capital from strategic competitors such as China and Russia. The Department of the Treasury has little insight into Russian and Chinese investments in U.S. emerging technology firms, as CFIUS filings are still largely voluntary for non-controlling investments in industries such as AI, semiconductors, quantum computing, and telecommunications equipment. While FIRRMA took positive steps in broadening CFIUS' authorities, it also left critical gaps in the investment screening regime. Additional steps are necessary to enable CFIUS to protect sensitive U.S. industries from adversarial capital, while ensuring the continued free flow of capital from trusted investors from allied nations.

Action for Congress:

- **Amend CFIUS' authorizing legislation to require competitors to disclose investments in “sensitive technologies” to CFIUS.**

- o Congress should amend CFIUS' authorizing legislation to mandate CFIUS filings for all non-controlling investments from “countries of special concern” in “sensitive technologies.” The Commission recommends that the legislation:
 - Define “countries of special concern” as states subject to export restrictions pursuant to section 744.21 of title 15 within the Code of Federal Regulations (China, Russia, and Venezuela) or any state that the Secretary of State designates as a state sponsor of terrorism (Iran, North Korea, and Syria).¹³
 - Require the Treasury Department to define a separate list of “sensitive technologies” for the purposes of CFIUS. Only investors from “countries of

special concern” would be required to submit CFIUS filings for investments in “sensitive technologies.” Treasury currently lacks authorities to broaden CFIUS’ mandatory filing requirements, which are linked to lists of technologies that are export controlled.¹⁴

- o Mandating CFIUS filings from select competitors in a broader set of sensitive industries—such as national security—relevant applications of AI, semiconductors, quantum computing, and advanced telecommunications equipment—will provide the Treasury with better visibility into Russian and Chinese investments in U.S. firms in key sectors. This allows CFIUS to operate with more precision and insight and focus attention on the riskiest investments.
- o Additionally, de-linking CFIUS disclosure requirements from export controls recognizes that there are instances in which it may be appropriate to screen investments prior to enacting export controls.¹⁵ Without this change, the only way to increase such disclosure requirements would be to place export controls on entire industries, which would significantly hamper commerce.

Action for the Department of the Treasury:

- **Expedite CFIUS exemption standards for allies and partners and create fast tracks for exempting trusted investors.**

- o The Department of the Treasury should issue clear guidance regarding what investment screening policies allied nations must implement to achieve CFIUS-exempted status.¹⁶ Clearly defining the standards for investment screening mechanisms in foreign nations necessary for investors to be exempted from CFIUS will create a powerful incentive for allied nations to adopt stronger screening mechanisms against adversarial capital. The sooner the Treasury takes this action, the more impact it will have on allied regulations. The Treasury should prioritize engagement with Five Eyes intelligence-sharing partners, Japan, South Korea, India, Israel, Singapore, Taiwan, and the European Union to enable investment from allied nations in U.S. high-tech firms.
- o Treasury should also issue new regulations creating a waiver for “trusted investors” from foreign countries that have a strong track record of CFIUS approval to exempt them from or lessen their CFIUS requirements. Currently there is no certification for investors with a trusted track record, and CFIUS treats foreign investors that are submitting for the first time the same as ones which have already submitted and been approved 100 times. Creating such a waiver would allow CFIUS to fast-track investments from low-risk, trusted investors with a strong history of CFIUS approval, facilitating legitimate foreign investment and focusing CFIUS’ resources on higher-risk investments.

Recommendation: Utilize Targeted Export Controls on Key Semiconductor Manufacturing Equipment

Recommendation

Although the Commission believes that export controls on AI algorithms would likely be ineffective given their widespread availability and commercial use, export controls on specific hardware components are capable of constraining competitors’ AI capabilities with national security applications and slowing their advancement. Policymakers must

be judicious in their application of such controls, as sweeping controls on general-use semiconductors are likely to cause substantial damage to the U.S. semiconductor industry and could have a net negative effect on overall U.S. competitiveness in microelectronics. However, targeted controls on key components that only the United States—or the United States and a small group of close allies—produce which are essential for cutting-edge defense applications could have a significant strategic impact at a relatively minimal cost.

The primary target for such controls should be select, high-end semiconductor manufacturing equipment (SME) needed to produce high-end chipsets, particularly photolithography equipment.¹⁷ China is the world's largest importer of SME, accounting for 29% of global imports from 2014 to 2018, and none of the largest or most sophisticated SME manufacturing firms are located in China.¹⁸ Simultaneous to implementing such controls, as discussed in Chapter 13 of this report, the United States should also fund efforts to prioritize the domestic development and manufacturing of SME tools and components needed to produce chips at scale at the 3nm node and beyond.¹⁹

Action for the Departments of Commerce and State:

- **Align the export control policies of the United States, the Netherlands, and Japan to restrict the export of high-end SME to China, including EUV and ArF immersion lithography equipment.**²⁰
 - o The Departments of State and Commerce should work to align the export control policies of the United States, the Netherlands, and Japan regarding high-end SME, particularly extreme ultraviolet lithography (EUV) equipment and argon fluoride (ArF) immersion lithography equipment, which is capable of producing chips at the 16nm node and below.²¹ All three states should establish a policy of presumptive denial of export licenses for exports of such equipment to China.²² This should include EUV scanner tools as well as specialized components for those tools, such as resist processing tools and EUV light sources, mirrors, and laser amplifiers. If such controls are effective, it will be difficult for China's government to cultivate indigenous, cutting-edge semiconductor fabrication capabilities and will degrade its advanced trailing-edge fabrication capabilities by complicating equipment repairs. Coupled with the refundable investment tax credit to promote U.S. semiconductor leadership recommended in Chapter 13 of this report, this will further the Commission's proposed U.S. policy goal of remaining two generations ahead of China in cutting-edge microelectronics design and fabrication.²³
- **Assess the effectiveness of existing U.S. export controls on SME on China's semiconductor industry and assess whether targeted controls on additional equipment are viable and necessary.**
 - o The Departments of Commerce and State should assess the effectiveness of existing U.S. export controls on SME on China's indigenous advanced semiconductor industry. Pending the results of that review and whether the Netherlands and Japan agree to align controls related to EUV and ArF immersion equipment, the United States could subsequently consider controls on additional SME chokepoints. If existing controls have failed to slow China's development of advanced fabrication capabilities, the United States could consider implementing controls on other targeted equipment chokepoints controlled by firms in allied

countries, such as atomic layer etching tools in conjunction with Japan and the United Kingdom.²⁴

Recommendation: Utilize End-Use Export Controls to Prevent Malicious Use of AI

Recommendation

Export controls that restrict transfer of dual-use items for specific end uses will not be effective at preventing technology transfer to determined adversaries, but they can still play a role in preventing the involvement of U.S. firms and technology in human rights abuses. For specific, high-end, dual-use equipment prone to facilitating uses of AI which enable human rights abuses, such as mass surveillance, U.S. firms should be required to certify that the equipment will not be used for specific nefarious ends and keep logs of their transactions. End-use controls and reporting requirements would not substantially delay sales and present a lower barrier to commerce compared to list-based controls. Requiring companies to self-certify and self-report could deter U.S. firms from knowingly enabling bad behavior abroad.

Action for the Department of Commerce:

- **Implement end-use controls and reporting requirements to prevent the use of high-end U.S. AI chips in human rights violations.**
 - o The Department of Commerce should implement end-use controls on high-end U.S.-designed or -manufactured AI chips for use in mass surveillance applications and institute reporting requirements on sales of such chips to China. The controls should be targeted only at very high-end or specialized chips, such as specific high-performing GPUs, ASICs, or FPGAs that exceed a certain high-performance threshold.²⁵ Commerce would, by necessity, update this threshold as chips continue to improve.
 - o Any firm that sells such chips to China should have to certify that the chips will not be used for any designated human rights abuses. Firms that sell such chips should also be required to provide quarterly reports to BIS listing all chip sales, in what quantity, and to which company. This will facilitate U.S. government tracking of chips that are most likely to facilitate abusive uses of AI and deter companies from selling chips to businesses that they know are engaging in such behavior.²⁶

Protecting the U.S. Research Environment

The United States needs comprehensive and resourced interagency measures to counter adversarial threats to its research environment, especially from China. Efforts must be supported by technically versed intelligence collection, analysis, and dissemination on threats in the Science & Technology (S&T) space. Promising steps have been initiated through the National Counterintelligence Task Force and the Office of Science and Technology Policy.²⁷ However, it is imperative to holistically improve the way the government postures itself and equips the research community—in academia and the private sector—to counter threats and uphold the integrity of open research.

Recommendation

Recommendation: Build Capacity to Protect the Integrity of the U.S. Research Environment

Actions for Congress:

- **Pass a modified version of the Academic Research Protection Act.**²⁸
 - o Congress should pass the Academic Research Protection Act (ARPA) with a modification that would mandate and execute standardization of grant processes across federal research–funding agencies.²⁹
 - The ARPA would establish a National Commission on Research Protection; establish an open–source intelligence clearinghouse relating to foreign threats to academia overseen by the Director of National Intelligence; improve guidance from the Departments of State and Commerce on export control responsibilities; and develop a Federal Bureau of Investigation (FBI) outreach strategy to promote information sharing on threats to the academic community.
 - The proposed modification would mandate development and implementation of a uniform application process and database across all Executive agencies that award R&D grants. This would enable effective oversight by grant–awarding agencies, allow for automated auditing, and support investigative efforts by federal law enforcement.
- **Establish a government–sponsored independent entity focused on research integrity.**
 - o Congress should authorize the sponsorship of a university–affiliated research center (UARC) to act as a center of excellence on research integrity and provide information and advice on research security.
 - o The entity should bridge the gap between the government and academic and private–sector research institutions and lower the barriers for research organizations to independently conduct compliance and informed risk assessments.
 - o The UARC mandate should be to:
 - Maintain open–source materials to serve university vetting of international engagement and risk management, including databases and risk–assessment tools;
 - Provide tailored guidance to research organizations for decision support;
 - Conduct comprehensive studies and regular reports on the state of foreign influence on U.S. research;
 - Undertake independent investigations on research integrity;
 - Develop education materials and tools for U.S. research institutions to build annual training and compliance initiatives; and
 - Manage dialogue with stakeholder communities and provide a venue for information sharing.

Action for the Director of National Intelligence:

- **Strengthen channels for information sharing with the research community.**
 - o In concert with the open-source intelligence clearinghouse relating to foreign threats to academia directed by the ARPA legislation, the Director of National Intelligence should support increased information and intelligence sharing with designated personnel at research organizations to share actionable information on specific threats. This would provide organizations the ability to swiftly take steps to mitigate risks.

Recommendation: Coordinate Research Protection Efforts Internationally with Allies and Partners

Recommendation

The United States should build a coalition of like-minded nations committed to the principle of open fundamental research and the associated values of research integrity—sidelining nations and organizations that do not abide by the values that provide the foundation for international innovation and science cooperation.³⁰

Action for the Office of Science and Technology Policy:

- **Foster international dialogue around research protection and integrity.**
 - o The Office of Science and Technology Policy, through the National Science and Technology Council, should work in coordination with Department of State's Office of Science and Technology Cooperation and Office of the Science and Technology Adviser to foster discussions with like-minded allies and partners focused on mitigating detrimental academic collaboration with China's People's Liberation Army (PLA)-affiliated and other high-risk entities. This should involve the establishment of an annual meeting of relevant education, science, and industry ministers to deepen research collaboration and coordinate on issues related to intellectual property and research security.

Action for the Department of Justice:

- **Strengthen information-sharing venues.**
 - o The Department of Justice (DOJ) and FBI, in coordination with Intelligence Community partners, should strengthen channels for information sharing on threats and best practices on research protection and coordinate multilateral responses to enforce research security.

Action for the Department of State:

- **Reinforce global norms around a commitment to open fundamental research.**
 - o Through international dialogues on research security and associated diplomacy, the Department of State should reinforce global norms around commitment to open

fundamental research,³¹ as described in the United States in the National Security Decision Directive (NSDD)-189, the *National Policy on the Transfer of Scientific, Technical and Engineering Information*.³²

Recommendation

Recommendation: Bolster Cybersecurity Support to Research Institutions

Protection of research data and intellectual property from cyber-enabled theft is perhaps the most important and actionable layer of security for the U.S. R&D environment. This is particularly true for AI, when theft of training data or trained models essentially provides malicious actors access to a final product. Federal investments in priority emerging technology research areas such as AI should be accompanied by a requirement and support for institutions—whether academic or private sector—to implement cybersecurity measures that adequately guard research data from cyber-enabled theft.

Actions for U.S. grant-making agencies:

- **Incentivize cybersecurity standards and best practices for grant-receiving research institutions.**
 - o U.S. grant-making agencies should provide incentives to research institutions to ensure that necessary practices, based on the existing NIST cybersecurity framework,³³ as well as governance processes are in place to protect sensitive research data.
 - Reporting structures and information flows of research institutions should be structured to raise cybersecurity as a critical issue for senior management and facilitate internal checks and audits. This includes senior leadership awareness of cyber threats, risk assessments, and active preventive measures.
 - U.S. grant-making agencies should make available incentives for research institutions that demonstrate adherence to cybersecurity standards and best practices.
 - Universities, research institutions, and other recipients of federal research funding should be required to periodically demonstrate that they are adhering to cybersecurity best practices. For government-owned and -sponsored laboratories, adherence to best practices, such as requiring critical data to be encrypted at rest and in transit, should be mandated and audited on a routine basis.
- **Support increased information sharing.**
 - o Research institutions receiving federal research dollars that do not already participate should be encouraged to join the Research and Education Networks Information and Sharing Analysis Center (REN-ISAC)³⁴ or an alternate ISAC, through which they can share information on threats and mitigation, benefit from automated threat-sharing tools, and have access to peer-assessment services to strengthen security postures.
 - o Similarly, research institutions should be made aware and encouraged to take advantage of the cybersecurity services offered by the Department of Homeland Security's Cybersecurity and Infrastructure Security Agency (CISA), to include automated indicator sharing³⁵ and enhanced cybersecurity services.³⁶

Action for the Federal Bureau of Investigation:

- **Share real-time, actionable threat information with research institutions.**
 - The FBI Cybersecurity Division should work closely with and share timely, anonymized threat information with REN-ISAC and research institutions to help them take active measures to counter cyber attacks and mitigate vulnerabilities.

Action for the Department of Homeland Security:

- **Support research cybersecurity information sharing similar to that of critical infrastructure.**
 - The Department of Homeland Security, CISA, and National Cybersecurity and Communications Integration Center³⁷ should support the level of information sharing with research institutions as they do with critical infrastructure and the Financial Services ISAC.³⁸

Action for the Office of Science and Technology Policy:

- **Support secure data storage.**
 - OSTP should broker commercial cloud credits³⁹ for universities to establish an ability to support secure data storage for research groups and laboratories conducting work known to be of high interest to foreign adversaries. This would provide an ability for universities to protect their sensitive research in a manner that does not require a significant capital investment.

Recommendation: Counter Foreign Talent-Recruitment Programs

Recommendation

China uses foreign talent-recruitment programs to achieve a “high ground” of AI experts.⁴⁰ Rather than pursue legitimate competition for scientific talent through attractive job offers, China’s talent-recruitment plans are designed in a manner that contradicts U.S. norms of research integrity, violates rules around disclosure, and creates vectors for technology transfer.⁴¹ The FBI and Intelligence Community assess that “participants are often incentivized to transfer to China the research they conduct in the United States, as well as other proprietary information to which they can gain access.”⁴² There is an urgent need to reinforce standards around disclosure of conflicts of interest and commitment and to create mechanisms that enable a heightened level of transparency and accountability.⁴³ This applies to researchers’ individual transparency and institutional accountability, as well as to the government in identifying problematic affiliations and enforcing standards. Currently, U.S. grant-making agencies lack common processes, coordination, and compliance mechanisms to enable this level of transparency and effective oversight.⁴⁴

Action for the Office of Science and Technology Policy:

- **Standardize grant application and recording processes.**
 - o The Office of Science and Technology Policy (OSTP), in coordination with the Office of Management and Budget, should provide advice and coordination to the Executive Branch to make uniform the grant application and recording processes across Federal agencies that fund external research.
 - o OSTP should advise and coordinate with agencies to ensure agencies embrace a government-wide standard for grant proposal documentation, requiring machine-readable formats that facilitate automation to identify fraud.⁴⁵ This would enable effective oversight by grant-awarding agencies, allow for automated auditing, and support investigative efforts by federal law enforcement.

Actions for Congress:

- **Mandate and resource compliance operations.**
 - o Congress should require and resource U.S. grant-making agencies to maintain compliance operations that can enforce standardized disclosure and accountability measures. Through periodic vetting and monitoring, grant-making agencies can provide a layer of accountability to enforce disclosure and protection policies.⁴⁶
- **Amend the Foreign Agent Registration Act.**
 - o Congress should amend the Foreign Agent Registration Act (FARA)⁴⁷ to require any individual or entity involved in the recruitment of U.S. nationals for a foreign talent program⁴⁸ to register as a foreign agent. This requires Congress to add a new category of activity to the legislation.

Actions for Department of Justice:

- **Update filing regulations to support an amended FARA.**
 - o Should Congress amend FARA legislation as proposed above, DOJ, in its implementing regulations, should identify specific information required from individuals involved in recruitment for foreign talent programs to ensure that the U.S. government has adequate visibility into foreign countries' talent recruitment activities in the United States.
 - o DOJ regulations should include methods for individuals and organizations to appeal a determination that they are subject to registration under this FARA expansion.
- **Publicly identify U.S.-based entities and foreign government proxies that serve as recruitment networks, platforms, or brokers.**
 - o To help raise awareness among researchers and research institutions, and reinforce transparency, Federal law enforcement and other relevant agencies should identify entities involved in recruitment activities for foreign talent programs and require their registration through the FARA (if amended).
 - o This effort must be accompanied by an associated appeal process for organizations to contest the need to register from identification.

Recommendation: Limit Collaboration with PLA-Affiliated Persons and Entities

PLA-affiliated universities and research labs send personnel abroad, with the overarching aim to obtain knowledge that can directly feed defense research and development priorities. Visiting scholars or students from PLA institutions often downplay their ties to the military or deliberately obscure affiliation by using alternate, external names for their home institutions that do not mention military or defense mandates.⁴⁹

The government should take actions through designation of institutions of concern and heightened visa vetting to assist universities in making risk assessments around research collaborations—becoming an effective partner in protecting research integrity.

Action for the Director of National Intelligence:

- **Create an open-source database of organizations that have a history of improper technology transfer, intellectual property theft, or cyber espionage.**⁵⁰
 - o The Director of National Intelligence, in coordination with law enforcement partners, should create a queryable database of academic institutions and other organizations that have a history of improper technology transfer, intellectual property theft, or cyber espionage. This resource should serve the research community and inform risk assessments of research organizations when entering collaborative arrangements. It would represent an expansive, open-source view of research institutions of concern, countering efforts to obscure military affiliations through adoption of innocuous institutional aliases.
 - o This must be accompanied by an associated appeal process for organizations to contest their inclusion in the database.

Action for the President:

- **Limit entrance of researchers with military and intelligence affiliations from countries of concern.**
 - o The President should issue an order to the Secretary of State and Secretary of Homeland Security to implement a requirement for special review of visas for advanced-degree students and researchers with ties to research institutions affiliated with foreign military and intelligence organizations of designated countries of concern.⁵¹
 - This should be paired with penalties that ban entry to any visa applicants found to have intentionally obscured institutional affiliations.

Action for the Department of State:

- **Resource special review measures.**
 - o Consular officers should be provided with adequate training, reference resources, analytical support, and time to conduct the special review.

Blueprint for Action: Chapter 14 - Endnotes

¹ A draft text of such an Executive Order is included in an Annex to this Blueprint for Action. This Executive Order also includes directives pertaining to most other export control–related recommendations in this Blueprint for Action.

² Additional details for this recommendation are also contained within the draft Executive Order included as an Annex to this Blueprint for Action.

³ Additional details for this recommendation are also contained within the draft Executive Order included as an Annex to this Blueprint for Action.

⁴ The ETTAC contains roughly 20 leading technical experts from prominent U.S. technology and defense firms, universities, and think tanks. However, it has been underutilized by Commerce; ETTAC did not hold a single meeting between June 2018 and May 2020. *Emerging Technology Advisory Committee*, Bureau of Industry and Security (last accessed Jan. 2, 2021), <https://tac.bis.doc.gov/index.php/ettac-home>.

⁵ Additional details for this recommendation are also contained within the draft Executive Order included as an Annex to this Blueprint for Action.

⁶ The Wassenaar Arrangement, a multilateral body with 42 participating states, is the primary international forum responsible for aligning policies on dual-use export controls. However, because it operates by consensus and includes Russia, is slow to react to new technologies and developments, and is non-binding, the Wassenaar Agreement must not be the exclusive forum in which the United States and allies negotiate export control provisions on dual-use technologies. *About Us, The Wassenaar Arrangement* (last accessed Jan. 2, 2021), <https://www.wassenaar.org/about-us/>; *Second Quarter Recommendations*, NSCAI at 68-69 (2020), <https://www.nscai.gov/previous-reports/>.

⁷ The Chapter 15 Blueprint for Action and associated Annex reinforce this recommendation and illustrate how these efforts should fit into a broader technology diplomacy strategy.

⁸ See Chris Darby, et al., *Mitigating Economic Impacts of the COVID-19 Pandemic and Preserving U.S. Strategic Competitiveness in Artificial Intelligence*, NSCAI at 14-15 (May 19, 2020), <https://www.nscai.gov/white-papers/covid-19-white-papers/>; *Second Quarter Recommendations*, NSCAI at 69, 75-77 (July 2020), <https://www.nscai.gov/previous-reports/>.

⁹ 50 U.S.C. § 4817(a)(1)(A).

¹⁰ 50 U.S.C. § 4565(a)(4)(B)(iii)(II); 85 Fed. Reg. 3112, *Provisions Pertaining to Certain Investments in the United States by Foreign Persons*, U.S. Department of Treasury: Office of Investment Security (Jan. 17, 2020) <https://www.federalregister.gov/documents/2020/01/17/2020-00188/provisions-pertaining-to-certain-investments-in-the-united-states-by-foreign-persons>.

¹¹ *New Controls on Emerging Technologies Released, While U.S. Commerce Department Comes Under Fire for Delay*, Gibson Dunn (Oct. 27, 2020), <https://www.gibsondunn.com/new-controls-on-emerging-technologies-released-while-us-commerce-department-comes-under-fire-for-delay/>; Letter from U.S. Senators Tom Cotton and Charles E. Schumer to Secretary Wilbur Ross, Department of Commerce (Nov. 18, 2019), https://www.cotton.senate.gov/imo/media/doc/191118_Cotton_Schumer_ECRA%20Letter%20to%20Sec.%20Ross%20copy.pdf.

¹² Additional implementation details for this recommendation are also contained within the draft Executive Order included as an Annex to this chapter.

¹³ *State Sponsors of Terrorism*, U.S. Department of State (last accessed Jan. 2, 2021), <https://www.state.gov/state-sponsors-of-terrorism/>.

¹⁴ As discussed in the following recommendation, due to the Department of Commerce's delay in identifying export controls on "emerging and foundational technologies," as required under the Export Control Reform Act of 2018 (ECRA), CFIUS' mandatory filing requirements have largely not expanded to emerging technology industries.

¹⁵ For instance, for early-stage technology venture investments, particularly those that do not yet produce specific products, export controls have historically been ineffective, but investment screening would still have value. See Michael Brown & Pavneet Singh, *China's Technology Transfer Strategy*, Defense Innovation Unit Experimental at 24 (Jan. 2018), [https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_\(1\).pdf](https://admin.govexec.com/media/diux_chinatechnologytransferstudy_jan_2018_(1).pdf).

¹⁶ CFIUS regulations released in January 2020 created an exception for non-controlling technology, infrastructure, and data (TID) investments for investors tied to “excepted foreign states,” with Australia, Canada, and the United Kingdom forming the initial list. The regulations require that excepted foreign states implement their own process to analyze foreign investments for national security risks and to facilitate coordination with the United States on investment screening by February 2022. However, Treasury has yet to publish the criteria CFIUS will use when determining whether additional countries can qualify as “excepted foreign states” in the future. See 31 C.F.R. 800.218 (2020), <https://home.treasury.gov/system/files/206/Part-800-Final-Rule-Jan-17-2020.pdf>.

¹⁷ The detailed reasons why high-end SME and photolithography equipment in particular represents the best target for such controls are described in Chapter 14 of this report.

¹⁸ John Verwey, *The Health and Competitiveness of the U.S. Semiconductor Manufacturing Equipment Industry*, SSRN at 5, 8 (July 1, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3413951.

¹⁹ See Chapter 13 of this report and its associated Blueprint for Action for additional details on recommendations to support the U.S. microelectronics industry, to include U.S. development of SME.

²⁰ Additional details for this recommendation are also contained within the draft Executive Order included as an Appendix to this chapter.

²¹ EUV lithography equipment is the only type of lithography equipment capable of mass manufacturing chips at the 5nm node or potentially below. ArF immersion lithography equipment is the only other type of tool capable of mass producing chips at the 28nm node or below, with more sophisticated ArF immersion equipment capable of nodes under 16nm. See Saif Khan, *Securing Semiconductor Supply Chains*, Georgetown Center for Security and Emerging Technologies at 20 (Jan. 2021), <https://cset.georgetown.edu/research/securing-semiconductor-supply-chains/>.

²² In 2019, the United States put significant pressure on the Netherlands to block a sale of EUV lithography equipment from Dutch firm ASML to Chinese firm SMIC. The contract expired before the equipment was delivered, although the Netherlands has not stated whether or not it will approve future sales. See Alexandra Alper, et al., *Trump Administration Pressed Dutch Hard to Cancel China Chip-Equipment Sale: Sources*, Reuters (Jan. 6, 2020), <https://www.reuters.com/article/us-asml-holding-usa-china-insight/trump-administration-pressed-dutch-hard-to-cancel-china-chip-equipment-sale-sources-idUSKBN1Z50HN>.

²³ Increasing the competitiveness of the cutting-edge U.S. microelectronics fabrication industry would create new market opportunities for SME firms, which could offset any potential losses resulting from decreased access to the Chinese market due to export controls. This is particularly important for allied governments that may be hesitant to impose export controls on equipment which will hurt key domestic companies without simultaneously providing them access to new markets or growing existing markets.

²⁴ Saif Khan, *Securing Semiconductor Supply Chains*, Georgetown Center for Security and Emerging Technologies at 20 (Jan. 2021), <https://cset.georgetown.edu/research/securing-semiconductor-supply-chains/>.

²⁵ GPUs are graphics processing units, ASICs are application-specific integrated circuits, and FPGAs are field-programmable gate arrays.

²⁶ The Chapter 15 Blueprint for Action reinforces this recommendation and illustrates how these efforts should fit into a broader technology diplomacy strategy.

²⁷ Specifically, the Joint Committee on Research Environments within the National Science and Technology Council. See *NSTC*, The White House (last accessed Jan. 1, 2021), <https://www.whitehouse.gov/ostp/nstc/>.

²⁸ H.R. 8346, Academic Research Protection Act, 116th Cong. (2020), <https://www.congress.gov/bill/116th-congress/house-bill/8346>.

Blueprint for Action: Chapter 14 - Endnotes

²⁹ This could mirror a provision for development of a uniform grant application process across research-funding agencies proposed in S. 3997, Safeguarding American Innovation Act, 116th Cong. (2020), <https://www.congress.gov/bill/116th-congress/senate-bill/3997/text>.

³⁰ Notably, two-thirds of overseas professional associations that transfer technology to China are located outside the United States. See Ryan Fedasiuk & Emily Weinstein, *Overseas Professionals and Technology Transfer to China*, Center for Security and Emerging Technology at 2 (July 21, 2020), <https://cset.georgetown.edu/research/overseas-professionals-and-technology-transfer-to-china/>. One-third of Thousand Talents awardees are located outside the United States, mainly in the U.K., Germany, and Singapore. See Ryan Fedasiuk & Jacob Feldgoise, *The Youth Thousand Talents Plan and China's Military*, Center for Security and Emerging Technology at 4 (Aug. 2020), <https://cset.georgetown.edu/research/the-youth-thousand-talents-plan-and-chinas-military/>. Two-thirds of awardees for some of China's largest scholarship programs are outside the United States. See Andrew Imbrie & Ryan Fedasiuk, *Untangling the Web: Why the US Needs Allies to Defend Against Chinese Technology Transfer*, Brookings Institution at 3 (April 2020), <https://www.brookings.edu/research/untangling-the-web-why-the-us-needs-allies-to-defend-against-chinese-technology-transfer/>. Leaders in Canada, the Netherlands, U.K., Japan, and India have in recent years publicly raised concerns around security risks related to research collaborations with China. Remco Zwetsloot, *China's Approach to Tech Talent Competition: Policies, Results, and the Developing Global Response*, Center for Security and Emerging Technology at 8 (April 2020), <https://cset.georgetown.edu/research/chinas-approach-to-tech-talent-competition-policies-results-and-the-developing-global-response/>.

³¹ This could build on a concept currently under consideration by the National Science Foundation to establish and formalize an international code of conduct around shared principles in research integrity and then fund collaborative research in accordance with said principles.

³² The directive defines fundamental research as: "'Fundamental research' means basic and applied research in science and engineering, the results of which ordinarily are published and shared broadly within the scientific community, as distinguished from proprietary research and from industrial development, design, production, and product utilization, the results of which ordinarily are restricted for proprietary or national security reasons." The key provision of NSDD-189 remains today: "It is the policy of this Administration that, to the maximum extent possible, the products of fundamental research remain unrestricted. It is also the policy of this Administration that, where the national security requires control, the mechanism for control of information generated during federally funded fundamental research in science, technology and engineering at colleges, universities and laboratories is classification." *National Policy on the Transfer of Scientific, Technical and Engineering Information*, NSDD-189 (Sept. 21, 1985), <https://fas.org/irp/offdocs/nsdd/nsdd-189.htm>.

³³ *Cybersecurity Framework*, NIST (last accessed Feb. 1, 2021), <https://www.nist.gov/cyberframework>.

³⁴ REN-ISAC (last accessed Jan. 2, 2021), <https://www.ren-isac.net/>.

³⁵ *Automated Indicator Sharing*, Cybersecurity and Infrastructure Security Agency (CISA) (last accessed Feb. 10, 2021), <https://www.cisa.gov/automated-indicator-sharing-ais>.

³⁶ *Enhanced Cybersecurity Services (ECS)*, Cybersecurity and Infrastructure Security Agency (last accessed Feb. 10, 2021), <https://www.cisa.gov/enhanced-cybersecurity-services-ecs>.

³⁷ *Cyber Incident Response*, CISA (Oct. 27, 2020), <https://www.cisa.gov/cyber-incident-response>.

³⁸ *Information Sharing and Awareness*, CISA (Dec. 8, 2020), <https://www.cisa.gov/information-sharing-and-awareness>.

³⁹ The National Science Foundation's CloudBank program could be leveraged as a model. See CloudBank, <https://www.cloudbank.org/>.

⁴⁰ William C. Hannes & Huey-meei Chang, *China's Access to Foreign AI Technology*, Center for Security and Emerging Technology (CSET) at 9-10 (Sept. 2019), https://cset.georgetown.edu/wp-content/uploads/CSET_China_Access_To_Foreign_AI_Technology.pdf.

⁴¹ The Office of Science and Technology Policy defines foreign government talent-recruitment programs as “an effort directly or indirectly organized, managed, or funded by a foreign government to recruit science and technology professionals or students (regardless of citizenship or national origin).” *Enhancing the Security and Integrity of America’s Research Enterprise*, Office of Science and Technology Policy at 18 (June 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/07/Enhancing-the-Security-and-Integrity-of-Americas-Research-Enterprise.pdf>.

⁴² Testimony of John Brown, Assistant Director, Counterintelligence Division, Federal Bureau of Investigation, delivered before the U.S Senate Committee on Homeland Security and Governmental Affairs, Permanent Subcommittee on Investigations, *Hearing on Securing the U.S. Research Enterprise from China’s Talent Recruitment Plans* at 2 (Nov. 19, 2019), <https://www.hsgac.senate.gov/imo/media/doc/Brown%20Testimony.pdf>. In some cases, the Chinese government appears to have rewarded scientists caught stealing technology through talent-recruitment programs, Alex Joske, *Hunting the Phoenix*, Australian Strategic Policy Institute at 8 (2020), <https://www.jstor.org/stable/resrep26119.1>.

⁴³ A National Science Foundation-commissioned JASON study on fundamental research security found that “disclosure of activities presents our main defense against foreign influence, especially that involving rewards, deception, and coercion.” *Fundamental Research Security*, JASON at 31 (Dec. 6, 2019), https://www.nsf.gov/news/special_reports/jasonsecurity/JSR-19-2IFundamentalResearchSecurity_12062019FINAL.pdf.

⁴⁴ *Threats to the U.S. Research Enterprise: China’s Talent Recruitment Plan*, U.S. Senate Permanent Subcommittee on Investigations (Nov. 2019), <https://www.hsgac.senate.gov/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China’s%20Talent%20Recruitment%20Plans.pdf>.

⁴⁵ This mirrors a recommendation from the U.S. Senate Permanent Subcommittee on Investigations. See *Threats to the U.S. Research Enterprise: China’s Talent Recruitment Plan*, U.S. Senate Permanent Subcommittee on Investigations at 11 (Nov. 2019), <https://www.hsgac.senate.gov/imo/media/doc/2019-11-18%20PSI%20Staff%20Report%20-%20China’s%20Talent%20Recruitment%20Plans.pdf>.

⁴⁶ The National Institutes of Health’s recent investments in this capability could serve as a model for others, scaled in terms of an agency’s level of funding.

⁴⁷ 22 U.S.C. § 611 et seq.

⁴⁸ This will require a clear definition of a foreign talent program, distinct from standard internationally funded research opportunities. The Office of Science and Technology Policy defines foreign government talent recruitment programs as “an effort directly or indirectly organized, managed, or funded by a foreign government to recruit science and technology professionals or students (regardless of citizenship or national origin).” *Enhancing the Security and Integrity of America’s Research Enterprise*, Office of Science and Technology Policy at 18 (June 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/07/Enhancing-the-Security-and-Integrity-of-Americas-Research-Enterprise.pdf>.

⁴⁹ Glenn Tiffert, *Global Engagement: Rethinking Risk in the Research Enterprise*, The Hoover Institution at 12 (2020), https://www.hoover.org/sites/default/files/research/docs/tiffert_globalengagement_full_0818.pdf.

⁵⁰ If Congress passes the Academic Research Protection Act, this initiative could be a component of the open-source intelligence clearinghouse on threats to academia created through the legislation.

⁵¹ This is recommended as an update to Presidential Proclamation 10043 that automatically suspends F or J visas to study or conduct research for Chinese nationals affiliated with the Chinese government military-civil fusion strategy. See Donald J. Trump, *Proclamation on the Suspension of Entry as Nonimmigrants of Certain Students and Researchers from the People’s Republic of China*, The White House (May 29, 2020), <https://trumpwhitehouse.archives.gov/presidential-actions/proclamation-suspension-entry-nonimmigrants-certain-students-researchers-peoples-republic-china/>. This order would provide for a case-by-case, risk-based review of potentially concerning applications from a broader group of designated countries.

Chapter 14 Annex: Technology Protection

Draft Executive Order on Export Control on Principles Guiding U.S. Policies for Protecting Dual-Use Technologies

By the authority vested in me as President by the Constitution and the laws of the United States of America, and in order to promote U.S. innovation and leadership in emerging and foundational technologies while protecting U.S. national security, it is hereby ordered as follows:

Section 1. Policy. It is the policy of the United States that export controls and investment screening mechanisms must be used in targeted, clearly defined, and strategic ways to protect U.S. national security, in pursuit of the broader policy of promoting U.S. innovation and leadership in emerging and foundational technologies, to include dual-use technologies such as artificial intelligence (AI).

The United States must be tailored and discrete in implementing export controls on dual-use emerging technologies such as AI. To ensure maximum effectiveness and minimize the adverse impact on U.S. industry, the U.S. Government should be guided by the following principles:

(1) **Principle One: Export Controls Cannot Supplant Investment and Innovation.** Technology protection policies are intended to slow U.S. competitors' pursuit and development of key strategic technologies for national security purposes, not stop them in their tracks. The United States must cultivate investment in these technologies through direct federal funding or changes to the regulatory environment in order to preserve existing U.S. advantages.

(2) **Principle Two: U.S. Strategies to Promote and Protect U.S. Technology Leadership Must Be Integrated.** The U.S. strategy to protect emerging technologies, including but not limited to AI, must be integrated with targeted efforts to promote U.S. leadership in such technologies. When choosing to implement controls, the United States should simultaneously consider policies to spur domestic research and development (R&D) in key industries to partially offset the resulting costs to U.S. firms, create alternative global markets, or encourage new investment to strengthen the U.S. industrial position.

(3) **Principle Three: Export Controls Must Be Targeted, Strategic, and Coordinated with Allies.** In devising new export controls on widespread and dual-use technologies such as AI, the United States must be careful and selective in the implementation of export controls. To ensure maximum effectiveness and minimize the adverse impact on U.S. industry, the U.S. Government should be guided by the following three-part test:

- a. Export controls must be targeted, clearly defined, and focused on choke points where they will have a strategic impact on the national security capabilities of competitors but smaller repercussions on U.S. industry.

b. Export controls must have a clear strategic objective, seeking to deter competitors from pursuing paths that endanger U.S. national security interests, and account for the projected cost and timeframe for competitors to create a domestic alternative.

c. Export controls must be coordinated with key U.S. allies which are also capable of producing the given technology, in order to effectively restrict the supply to adversaries and also prevent circumstances in which unilateral controls cut off U.S. market access but competitors are able to purchase the same technology from other countries.

(4) Principle Four: The United States will be judicious in its use of export controls but broaden investment screening on critical and emerging technologies.

While broad and sweeping export controls on AI and other dual-use emerging technologies could result in significant blowback on U.S. industry, which would harm overall U.S. strategic competitiveness, investment screening presents opportunities to take a more proactive regulatory approach while minimizing risk to U.S. industry. Provided the United States can continue approving benign transactions expeditiously, enhancing investment screening presents significant potential to blunt concerning transfers of technology.

Section 2. Objective. In 2018, the Congress enacted the Export Control Reform Act of 2018 (ECRA) and the Foreign Investment Risk Reduction Modernization Act of 2018 (FIRRMA) to provide the U.S. Government with additional mechanisms to control exports and screen investments. The U.S. Government must take steps to provide the private sector and foreign governments with clarity about the application of these laws to emerging and foundational technologies and enhance U.S. national security in the process.

Section 3. Establishment of Interagency Task Force on Emerging and Foundational Technologies. (a) Pursuant to Section 1758 of ECRA, there is hereby established an Interagency Task Force on Emerging and Foundational Technologies (Task Force) to identify emerging and foundational technologies that are essential to the national security of the United States and are not critical technologies described in clauses (i) through (v) of 50 U.S.C. 4565(a)(6)(A).

(b) The Task Force shall be chaired by the Secretary of Commerce (Chair) and consist of senior-level officials from the following Executive departments and agencies (agencies) designated by the heads of those agencies:

(i) Department of State;

(ii) Department of the Treasury;

(iii) Department of Defense;

(iv) Department of Energy; and

(vi) such other agencies as the President, or the Chair, may designate.

(c) The Chair shall designate a senior-level official of the Department of Commerce as the Executive Director of the Task Force, who shall be responsible for regularly convening and presiding over the meetings of the Task Force, determining its agenda, and guiding its work in fulfilling its functions under this Order, in coordination with the BIS at the Department of Commerce.

Section 4. Functions of the Task Force.

(a) The Task Force shall meet regularly to identify emerging and foundational technologies that are essential to the national security of the United States for purposes of establishing export controls and investment screening mechanisms, as appropriate, related to those technologies.

(b) Within 120 days, the Task Force shall finalize lists of emerging and foundational technologies pursuant to section 1758 of ECRA. The Secretary of Commerce shall thereafter issue proposed rules on emerging and foundational technologies and proceed expeditiously to issue final rules at the conclusion of the notice and comment period.

(c) The Task Force shall review the lists of emerging and foundational technologies and issue amendments as needed on no less than an annual basis.

Section 5. Process for Identifying Emerging and Foundational Technologies.

(a) In identifying emerging and foundational technologies pursuant to this Order, the Task Force shall consider information from multiple sources, including:

(i) publicly available information;

(ii) classified information, including relevant information provided by the Director of National Intelligence;

(iii) information relating to reviews and investigations of transactions by the Committee on Foreign Investment in the United States under 50 U.S.C. 4565; and

(iv) information provided by the advisory committees established by the Secretary to advise the Under Secretary of Commerce for Industry and Security on controls under the Export Administration Regulations, including the Emerging Technology Technical Advisory Committee (ETTAC).

(b) In identifying emerging and foundational technologies pursuant to this Order,

the Task Force shall take into account:

- (i) the development of emerging and foundational technologies in foreign countries;
- (ii) the effect that export controls imposed pursuant to this section may have on the development of such technologies in the United States;
- (iii) the effectiveness of export controls imposed pursuant to this section on limiting the proliferation of emerging and foundational technologies to foreign countries; and
- (iv) the policy and principles reflected in section 1 of this Order.

Section 6. Improving Coordination with Expert Advisory Groups. (a) The Secretary of Commerce shall review existing technical advisory committees (TACs) at the Department of Commerce, including the ETTAC, to ensure that each TAC is composed of members from industry and academia with deep subject-matter expertise to assess the need for export controls for emerging and foundational technologies.

(b) The Secretary of Commerce, as Chair of the Task Force, shall ensure that the Task Force has solicited and received feedback from the ETTAC and other relevant TACs at the Department of Commerce on the text of any proposed or final rule on emerging or foundational technologies, prior to issuance of such rule.

(c) The Secretary of Commerce shall ensure that senior officials at the Departments of State and the Treasury are granted non-voting observer access at all ETTAC meetings.

Section 7. Improving International Coordination on Export Controls on Semiconductor Manufacturing Equipment. Within 180 days, the Secretary of State, in consultation with the Secretary of Commerce and the Secretary of Defense, shall host a multilateral engagement with senior-level representatives of Japan, the Netherlands, and, if deemed appropriate, other U.S. allies and partners that produce semiconductor manufacturing equipment (SME), including EUV lithography equipment and ArF immersion lithography equipment, listed by the Wassenaar Arrangement or identified by the Task Force. The purpose of this meeting will be to align export licensing policies toward a presumptive denial of export licenses for exports of semiconductor manufacturing equipment to China. The Secretary of State shall provide a report to the President within 60 days of the meeting assessing:

- (i) whether U.S. allies and partners are currently exporting such equipment to China;

(ii) what steps each country that manufactures such equipment must take to ensure its regulatory regime is aligned with that of the United States, and its willingness to take those steps; and

(iii) whether additional opportunities exist to strengthen international cooperation on export controls on SME which are consistent with the policy and principles reflected in Section 1 of this Order.

Section 8. Engaging Technical Experts for Export Control Review. (a) The Secretary of Commerce, in consultation with the Secretaries of the Treasury and Defense, shall establish a network within existing Federally Funded Research and Development Centers (FFRDCs) and University Affiliated Research Centers (UARCs) to provide technical expertise to all departments and agencies for issues relating to export controls and investment screening related to emerging and foundational technologies. The network shall encompass a regional distribution of FFRDCs and UARCs located in areas of the United States with a concentration of technology expertise in emerging and foundational technologies.

(b) Individuals selected to participate in the network shall provide real-time technical input to all policy discussions on export controls and review of export control license applications, including those of the Task Force, those conducted pursuant to EO 12981 or a successor order, and any other interagency policy discussions pertaining to export controls, as well as the investment screening processes of the Committee on Foreign Investment in the United States (CFIUS).

Section 9. Automating Export Control and Investment Screening Reviews. The Secretaries of Commerce and the Treasury shall task the aforementioned network with exploring using AI-based systems to assist in the evaluation of applications for export control licenses and CFIUS filings and shall provide a report to the President on the use of AI-based systems for such purposes within 180 days. This report shall include an evaluation of:

(i) how AI-based systems could assist existing review processes;

(ii) whether incorporating such systems could enhance the accuracy and speed of the review processes;

(iii) whether relevant Departments and Agencies have sufficient quantity and quality of data to train AI-based review systems, and how existing data can be improved;

(iv) what information technology infrastructure inside relevant Departments and Agencies needs to be improved to fully utilize such systems; and

(iv) an approximate timeline and cost for deploying a system or systems, and the projected savings per year in labor-hours once deployed.

Section 10. General Provisions. (a) Nothing in this Order shall be construed to impair or otherwise affect:

(i) the authority granted by law, regulation, Executive Order, or Presidential Directive to an executive department, agency, or the head thereof; or

(ii) the functions of the Director of the Office of Management and Budget relating to budgetary, administrative, or legislative proposals.

(b) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.

(c) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.

Chapter 15: A Favorable International Technology Order

Blueprint for Action

This Blueprint for Action provides detail for a comprehensive strategy to further U.S. interests with allies and partners to shape a favorable international technology order, win the technology competition against authoritarian states, and advance artificial intelligence (AI) innovation and adoption across the world to promote the values of free and open societies. This Blueprint for Action also proposes reforms to reorient U.S. foreign policy and the Department of State for great power competition in the digital age.

Recommendation

Recommendation: Develop an International Science & Technology Strategy

The International Science & Technology Strategy (ISTS) will help coordinate emerging technology policies across the government and with our closest allies and partners; apply the tools of foreign assistance, technical expertise and guidance, and development finance and investment; and foster collaborative R&D. The ISTS should serve as the international component of the National Technology Strategy (NTS) and provide an organizing framework to drive U.S. foreign policy with regard to emerging technologies.¹ The ISTS should center on four big initiatives:

- Building an Emerging Technology Coalition (ETC);
- Launching an International Digital Democracy Initiative (IDDI);
- Implementing a coordinated U.S. national plan to support international efforts; and
- Enhancing the United States' position as an international digital research hub.

Action for the President:

- **Direct development of an International Science & Technology Strategy (ISTS) by a White House-led interagency task force.**
 - o The President should direct development of the ISTS by a dedicated task force.
 - o The ISTS Task Force should be convened by the Technology Competitiveness Council or otherwise co-chaired by the Assistant to the President for National

Security Affairs and the Directors of the Office of Science and Technology Policy and the National Economic Council.

- o The ISTS Task Force should include leadership from the following agencies:
 - o the Department of State;
 - o the Department of the Treasury;
 - o the Department of Commerce, including the Bureau of Industry and Security (BIS) and the National Institute of Standards and Technology (NIST);
 - o the Department of Energy (DOE);
 - o the Department of Homeland Security's Cybersecurity & Infrastructure Security Agency (CISA);
 - o the National Science Foundation (NSF);
 - o the United States Agency for International Development (USAID);
 - o the U.S. International Development Finance Corporation (DFC);
 - o the Export-Import Bank of the United States (EXIM);
 - o the U.S. Trade and Development Agency (USTDA);
 - o the Millennium Challenge Corporation (MCC); and
 - o as appropriate, other agencies with expertise on individual topics.
- o The ISTS Task Force should develop and submit to the President a formal strategy, linked closely to the President's National Security Strategy (NSS) and the Secretary of State's and USAID Administrator's Joint Strategic Plan (JSP), building on those documents' technology-related goals and priorities. The ISTS should serve as the international component of the National Technology Strategy.²
- o The ISTS should be centered around four big initiatives addressed in this Plan:
 - building an Emerging Technology Coalition;
 - launching an International Digital Democracy Initiative;
 - implementing a comprehensive U.S. national plan to support international digital efforts around technical standards, foreign assistance, development finance, and export controls; and
 - enhancing the United States' position as an international digital research hub.
- o Once approved by the President, the ISTS Task Force would be responsible for overseeing and supporting the implementation, to include identifying resource and organizational changes needed to implement the strategy. The ISTS Task Force should hold regular meetings to facilitate execution of the strategy.

Recommendation

Recommendation: Build an Emerging Technology Coalition

As part of the ISTS, the United States, led by the White House and the Department of State, should lead in forming an Emerging Technology Coalition (ETC) of countries respectful of democratic values. The ETC would be a body of like-minded allies and partners to work with each other and with help from international and non-governmental organizations, civil society actors, and the private sector to develop and implement a coordinated strategy and associated policies to:

1. promote the design, development, and use of emerging technologies according to democratic norms and values;
2. coordinate policies and investments to counter the malign use of these technologies by authoritarian regimes; and
3. provide concrete, competitive alternatives to counter the adoption of digital infrastructure made in China.

Action for the White House and the Department of State:

- **Convene key allies and partners to join and establish the ETC.**

- o The United States should lead an ETC of like-minded nations either as part of a larger democracy summit or as a stand-alone endeavor.
- o Membership should include a core group of technologically advanced democratic nations, reflecting a broad geographic distribution, that have demonstrated shared interests in advancing responsible AI, countering malign uses of emerging technologies, and ensuring high standards for openness, trust, and privacy in digital infrastructure.
 - The ETC should build on two important dialogues previously recommended by the Commission: the U.S.-India Strategic Tech Alliance and the U.S.-EU Strategic Dialogue for Emerging Technologies.³
 - The ETC should build on—and be additive to—promising efforts and projects underway at the Organization for Economic Co-Operation and Development (OECD) and the Global Partnership on AI (GPAI).⁴ *See Table 1. Key Multilateral Technology Initiatives (located at the end of this Plan).*
- o The Commission further recommends that the ETC invite representatives from international organizations, non-governmental organizations, civil society, academia, and the private sector.
 - These organizations are critical to implement policies across borders, convene state and non-state actors, and promote alignment of responsible AI and digital infrastructure development and use in accordance with shared democratic values.⁵
 - They should be included in the ETC, among participants in the inaugural session, and should have observer status.

Actions for the United States and Allies and Partners:

- **Organize efforts to synchronize policies and initiatives across seven critical areas.**
 - o The ETC should be organized around a concrete agenda with actionable objectives focused on the outcomes rather than processes, designed to develop and realize a shared vision of a positive technological future and contrast it against a future dominated by authoritarian practices.
 - o Building on an existing framework of guiding principles, such as the OECD AI Principles,⁶ members should use the inaugural meeting to endorse a concrete agenda designed to operationalize policies and initiatives across seven critical areas:
 - *Developing and operationalizing standards and norms* in support of democratic values and the development of secure, reliable, and trusted technologies;
 - *Promoting and facilitating coordinated and joint R&D on AI and digital infrastructure* that advances shared interests and benefits humanity;
 - *Promoting democracy, human rights, and the rule of law* through joint efforts to counter censorship, malign information operations, human trafficking, and illiberal uses of surveillance technologies;
 - *Exploring ways to facilitate data sharing* among allies and partners through enabling agreements, common data archival procedures, cooperative investments in privacy-enhancing technologies, and addressing legal and regulatory barriers;
 - *Promoting and protecting innovation*, particularly through export controls, investment screening, supply chain assurance, emerging technology investment, trade policy, research and cyber protections, and intellectual property alignment;
 - *Developing AI-related talent* by analyzing labor market challenges, harmonizing skills and certification requirements, and increasing talent exchanges, joint training, and workforce development initiatives; and
 - *Launching the International Digital Democracy Initiative* to coordinate international foreign assistance, development aid and financing, technical guidance, and policy guidance.
 - o To execute an agenda across the seven critical areas, the ETC members should consider creating implementation groups for each area.
 - o Proposed agendas and guidance for each critical area are included in the Emerging Technology Coalition Annex to this Blueprint for Action.

Recommendation: Launch an International Digital Democracy Initiative

Recommendation

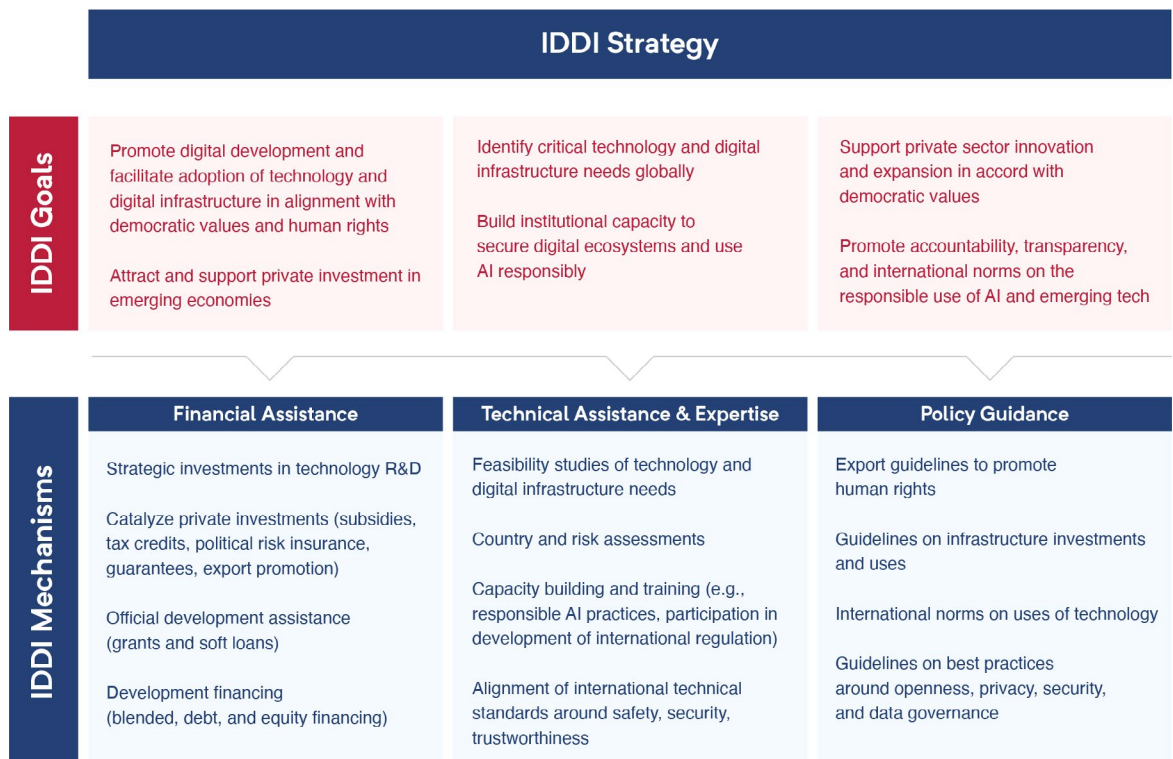
The Commission recommends that the United States and ETC partner states launch an International Digital Democracy Initiative (IDDI), a coordinated effort to align partner states' foreign assistance policies and programs to develop, promote, and fund the adoption of AI and associated technologies that comport with democratic values and ethical norms around openness, privacy, security, and reliability.⁷ The IDDI will:

- Coordinate partner-state approaches to adopting and governing digital technologies;
- Mobilize coalition efforts to provide alternatives (through funding assistance, technology development, and private-sector investment) to digital infrastructure and AI/machine language (ML)-enabled technologies that are used for illiberal ends and to promote technologies that enhance democratic participation, human rights, and the rule of law; and
- Facilitate adoption of secure, reliable, and trusted digital infrastructure, AI/ML-enabled technologies, and information and communications technology (ICT).⁸

Actions for the United States and Allies and Partners:

- **Coordinate national strategies that articulate involvement in IDDI.**
 - o The United States and IDDI partners should take steps to coordinate the development of national strategies for IDDI involvement. By focusing on developing and investing in democratically aligned digital technologies and supporting digital development, infrastructure, and capacity-building projects, national strategies for IDDI should further the overarching goals of the ETC. *The figure below provides an overview of the IDDI.*
 - o The Commission recommends that IDDI partners seek to align national strategies around common guidelines for investment strategies, critical technologies, policy guidance, and export promotion. A public diplomacy plan and associated resources should also be prominent within each national strategy given the importance of promoting a positive, unified message on the benefits and importance of IDDI.

Overview of IDDI Strategy.



- **Conduct an assessment of the global digital development landscape.**
 - o IDDI partner states should convene with representatives from development agencies and international financial institutions (IFIs) to conduct an assessment of digital connectivity and the global digital development environment to guide IDDI activities.⁹
 - o The Commission proposes that this assessment include:
 - A global risk evaluation of state-sponsored policies, financing and investment tools, surveillance technologies, and other mechanisms that erode privacy and civil and human rights. This evaluation would inform IDDI priorities.
 - Identification of technologies or technological features to promote through IDDI activities, incorporating some or all of the following:
 - privacy protections, such as privacy-preserving ML, eyes-off ML, advanced encryption, and secure multi-party computational models¹⁰;
 - protections against unwanted bias in data and inferences;
 - restrictions on the use of certain applications to prevent the potential infringement on civil and human rights;
 - data storage and access restrictions, to prevent access from third parties, multiple government agencies, and foreign governments;
 - secure, reliable AI tools and digital infrastructure;
 - tools and infrastructure to support “green” initiatives, including smart grids; and
 - tools for local populations to counter authoritarian, social-control uses of AI.
 - Identification of best practices within IDDI members and existing initiatives that provide solid foundations to build upon and develop at scale. The IDDI should capitalize on the unique capabilities and resources of individual IDDI partner states.
 - *See Table 2. International Digital Development Programs (located at the end of this Plan) and the figure below.*

Models for International Digital Democracy Initiative

The following examples of efforts to develop, promote, and fund the adoption of secure, trusted, and open digital ecosystems can serve as models for IDDI projects.

Models for International Digital Democracy Initiative (IDDI).

Digital Inclusion
USAID partners with the private sector to close the gender digital divide, particularly by increasing internet coverage of women in Colombia, Guatemala, Ghana, India, and Kenya.

Three Seas Initiative (3SI)
In Dec. 2020, 3SI announced its first digital infrastructure investment: Greenery Data Centers will develop data centers throughout the region.

Blue Dot Network
Launched in Nov. 2019, the U.S., Japan, and Australia are helping countries (particularly in the Indo-Pacific) through risk assessments and quality infrastructure investments that attract sustainable, transparent development.

Digital Connectivity and Cybersecurity Partnership (DCCP)
USAID has provided policy support, technical assistance, capacity-building, and trainings throughout Latin America and the Caribbean on ICT policy and regulation.

FAIR Forward — AI for All
The German initiative—in partnership with Ghana, Rwanda, South Africa, Uganda, and India—seeks an “open, inclusive and sustainable approach to AI” in Africa and Asia through training, improving access to data and AI technologies, developing policy, and supporting AI accelerators.

Trilateral Infrastructure Partnership (TIP)
In Oct. 2020, the U.S., Japan and Australia announced joint funding for a \$30M undersea fiber optic cable to the Republic of Palau.

Women’s Global Development and Prosperity (W-GDP) Fund Announces \$122 Million in Progress and Partnerships, USAID (Sept. 3, 2020), <https://www.usaid.gov/w-gdp/fact-sheet/aug-2020-womens-global-development-and-prosperity-fund-announces-122m-progress-partnerships>; *Digital Connectivity and Cybersecurity Partnership (DCCP)*, USAID (Oct. 19, 2020), <https://www.usaid.gov/digital-development/digital-connectivity-cybersecurity-partnership>; Press Release, U.S. Embassy Chile, *U.S. Support for Digital Transformation in Latin America and the Caribbean* (Nov. 10, 2020), <https://cl.usembassy.gov/u-s-support-for-digital-transformation-in-latin-america-and-the-caribbean/>; *The Three Seas Fund Makes Its First Digital Investment*, Three Seas Initiative Investment Fund (Dec. 2, 2020), <https://3siif.eu/news/the-three-seas-fund-makes-its-first-digital-investment>; *FAIR Forward - Artificial Intelligence for All*, Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH (June 2020), <https://toolkit-digitalisierung.de/app/uploads/2020/06/Factsheet-FAIR-Forward-E050620-1.pdf>; *The Launch of Multi-Stakeholder Blue Dot Network*, U.S. International Development Finance Corporation (Nov. 4, 2019), <https://www.dfc.gov/media/opic-press-releases/launch-multi-stakeholder-blue-dot-network>; *Blue Dot Network*, U.S. Department of State (last accessed Feb. 16, 2021), <https://www.state.gov/blue-dot-network/>; Fact Sheet, U.S. Department of State, *The United States Partners with Australia and Japan to Expand Reliable and Secure Digital Connectivity in Palau* (Oct. 29, 2020), <https://2017-2021.state.gov/the-united-states-partners-with-australia-and-japan-to-expand-reliable-and-secure-digital-connectivity-in-palau/index.html>; Press Release, Australian Infrastructure Financing Facility for the Pacific, *Australia Partnering with Japan and the United States to Finance Palau Undersea Cable* (Oct. 28, 2020), <https://www.aifpp.gov.au/news/australia-partnering-japan-and-united-states-finance-palau-undersea-cable#:~:text=This%20project%2C%20valued%20at%20approximately,Japan%20and%20the%20United%20States.>

- **Determine investment guidelines for technology development and digital development projects and support alignment through OECD.**
 - o Investment decisions into the development of specific technologies and funding of various digital development projects should be guided by the outcome of the assessment and agreed-upon digital development and AI use principles. These include:
 - G20 Principles for Quality Infrastructure Investment ¹¹;
 - Principles for Digital Development, used by USAID to guide digital foreign assistance efforts¹²;
 - Criteria for Security and Trust in Telecommunications Networks and Services, developed by the Department of State, the Center for Strategic and International Studies, and outside experts, and used by the Blue Dot Network¹³;
 - OECD's Recommendation on Digital Security of Critical Activities¹⁴;
 - the forthcoming OECD Principles on Trusted Government Access to Data;
 - OECD Guidelines for Multinational Enterprises¹⁵; and
 - UN Guiding Principles on Business and Human Rights.¹⁶
 - o IDDI members should also work with the OECD to standardize AI and digital development assistance through the creation of a dedicated "Digital Development" purpose code.
 - A dedicated purpose code, monitored by the OECD's Development Assistance Committee (DAC),¹⁷ will demonstrate international resolve, facilitate coordination, enable the OECD and other entities to monitor funding in digital development activities, and consolidate data to inform IDDI strategic decision-making.
- **Develop guidelines for the use of technologies within the IDDI.**
 - o The risk assessment should lead to the development of guidelines for the use of AI/ML-enabled applications and surveillance technologies.
 - This effort should build on several foundational documents, including the OECD AI Principles¹⁸ along with NSCAI's Key Considerations for Responsible Development & Fielding of Artificial Intelligence,¹⁹ which provide operational guidance for the responsible and ethical development and use of AI in engineering practices, system performance, human-AI interaction, and accountability and governance.
- **Develop export promotion and control principles and coordinate adoption by each partner state.**
 - o IDDI nations should establish priorities for export promotion and R&D activities to promote technologies that comport with shared democratic values and support free and open societies. These priorities may expand upon OECD guidelines and new U.S. Department of State guidelines on surveillance due diligence (see below on Promoting Democracy through Export Controls), to incentivize companies against transactions that could result in misuse of technology by government end users.²⁰

- **Expand public- and private-sector investments by exploring the creation of a joint investment fund and incentives for private investment.**
 - o IDDI members should consider creating a joint investment fund—with a dedicated investment manager—to support IDDI projects. Such a fund could be modeled on the Three Seas Initiative Investment Fund (3SIIF). *See Table 2. International Digital Development Programs (located at the end of this Plan).*
 - o As public-sector investment is unlikely to achieve the scale necessary to realize IDDI goals and safeguard IDDI partner states' collective security,²¹ IDDI partners should seek to catalyze at least \$20 billion in private-sector investment.
 - o IDDI members could explore incentives for private-sector investment in critical emerging technologies, particularly in the Indo-Pacific, Latin America, and other regions with strong growth potential. Policies to explore include tax incentives and subsidies, communication of IDDI priorities to the private sector, highlighting private-sector investments and practices that advance IDDI goals, and increased taxes on profits made from strategic competitors' publicly traded companies.²²
- **Execute a coordinated strategic messaging and awareness campaign.**
 - o The success of the IDDI will depend not only on coordinated investment and assistance activities, but also on the ability of IDDI members to effectively and strategically communicate the objectives to world leaders, international organizations, and the public.

Recommendation

Recommendation: Develop and Implement a Comprehensive U.S. National Plan to Support International Technology Efforts

The ISTS should include an integrated government-wide plan for coordinating the tools of U.S. foreign policy to advance the ETC, the IDDI, and stand-alone projects. This plan should leverage technical expertise, foreign assistance, development financing and investment, policy guidance, and export controls in support of three core goals:

1. Shaping international technical standards on AI and related technologies;
2. Implementing a coordinated U.S. policy for the IDDI; and
3. Promoting transparency and accountability through export controls.

Component 1: Shape International Technical Standards

The United States and its allies should lead the way on international technical standardization for AI. U.S. government-led dialogue with U.S. industry, as well as democratic allies, can help overcome information asymmetries and clarify objectives for technical standards on AI that foster economic growth, protect consumers, and safeguard democratic values. Partnership and information-sharing between the U.S. government, industry, and academia is critical to ensure protection of national security concerns involving standards and the neutrality of international standards-setting bodies.²³

Action for the President:

- **Issue an Executive Order to support international technical standardization.**

- o As detailed in NSCAI's *Interim Report and Third Quarter Recommendations*, the President should issue an executive order²⁴ that would:
 - establish an interagency coordination task force for sharing threat information and identifying U.S. national security interests related to AI technical standards, and related standards such as international data science standards, to be led by NIST with membership from the Departments of State, Defense, Energy, Commerce, and Homeland Security, the Office of the Director of National Intelligence, and USAID;
 - direct the interagency task force to improve partnership and collaboration with industry and academia;
 - direct the interagency task force to consult with relevant congressional committees and develop a work plan with congressional appropriators on the necessary resources and full-time equivalents necessary to support U.S. leadership in international technical standardization;
 - direct federal agencies to resource and support focused research, test, and evaluation and regular and active participation by the U.S. Government in international standards-setting activities;
 - require the Director of NIST and the Standards Coordinator to encourage the private sector to create a Standardization Center to improve sharing of best practices and other information relevant to standards development, as well as support focused research coordination; and
 - establish a federal advisory committee with experts from the private sector and academia to provide strategic guidance to the interagency coordination task force on international technical standards.

Action for the Department of Commerce:

- **Coordinate technical standards-development activities government-wide through NIST leadership of the interagency task force.**

- o The development of international standards for AI and emerging technology should be incorporated into the overarching ISTS. Within the U.S. government, this process must continue to be led by NIST with active participation of agencies in the coordination task force described above.
- o The Commission has proposed a comprehensive plan for NIST and other U.S. departments and agencies to ensure that the development of international technical standards receives greater attention and resourcing to ensure that U.S. national security interests, including the promotion of technologies that comport with democratic values, are advanced in standards-development organizations.²⁵
- o NIST and other agencies should consider the Commission's *Key Considerations for Responsible Development and Fielding of AI* in assessing positions on technical standards.²⁶

- **Convene a federal advisory committee to inform strategy on international standards.**
 - o As noted above, the proposed Executive Order would create a federal advisory committee to provide the interagency task force with expert guidance to inform U.S. government strategy on international technical standards.
 - o Members of the advisory committee should be drawn from the private sector and academia and should be selected by the interagency task force for their expertise in emerging technologies, geopolitical analysis, global economic trends, and similar fields.
 - o The Commission envisions that this advisory committee, by focusing on strategic geopolitical issues around international technical standards, would serve a function not currently fulfilled by other advisory groups and the industry organizations that coordinate U.S. positions before international standards bodies.²⁷
 - o The advisory committee should have a forward-looking mandate to contribute to U.S. government strategy on a range of emerging technologies—including technologies involved in genomics, digital currency, biopharma production, and others.
 - o NIST and the Department of State should ensure that members receive appropriate clearances to facilitate exchanges of classified information necessary to the development of U.S. strategy.

Action for the Departments of Commerce and State:

- **NIST, with assistance from the Department of State, should coordinate technical standards-development activities internationally.**
 - o In addition, NIST, working closely with the Department of State—ideally, in the context of the ETC and the IDDI—must prioritize engagement with democratic nations to align positions on standards critical to mutual security and defense and ensure those positions are reflected in deliberations of technical standards-development organizations.
 - o The Department of State's Regional Technology Officers can serve as conduits for this alignment (see below on "Reorient U.S. Foreign Policy and the Department of State").

Actions for Congress:

- **Provide appropriate funding to NIST and other U.S. departments and agencies to support international technical standardization efforts.**
 - o As the Commission has recommended,²⁸ Congress should provide funds sufficient to support at least six full-time equivalent personnel at NIST and at least one full-time equivalent each at the Departments of State, Defense, Homeland Security, and Energy; the Office of the Director of National Intelligence; USAID; and other agencies as may be appropriate. These personnel will support NIST's AI Standards Coordinator, support focused research, and undertake other responsibilities

necessary for technical standardization, such as participating in standards-development organizations.

- **Provide appropriate funding, and grant-issuing authority, for the Department of State to ensure international leadership in developing technical standards.**
 - o As the Commission has recommended,²⁹ the Department of State must be properly resourced to fully engage in international forums, unions, and organizations focused on developing standards for AI, associated technologies, and data. Congress should provide a minimum of \$5 million to support these endeavors, particularly the recruitment and funding of U.S. academic scholars and researchers to participate in these international forums. This action may require the creation of a new foreign assistance fund and grant-issuing authority to a Department office.
- **Establish a grant program to enable small- and medium-sized U.S. AI companies to participate in international standardization efforts.**
 - o As the Commission has recommended,³⁰ Congress should authorize a grant program for small- and medium-sized U.S. AI companies to cover the high costs of engaging in international standardization efforts, including conducting relevant research, developing requisite skills and expertise, preparing standards proposals, and attending technical standards-setting meetings. Their input enables greater technological innovation, helps prevent potential high “switching costs” that may impede their growth, and facilitates solution development for standards that impede exporting by these small businesses.
 - o The Commission proposes that Congress appropriate an initial amount of \$1 million annually to fund grants issued by the Small Business Administration, in coordination with NIST.

Component 2: Implement a Coordinated U.S. National Policy for the IDDI

A national policy for U.S. digital development efforts and involvement in IDDI will provide high-level strategic vision and coordination necessary to:

- Advance the interests of the United States and its allies and partners in the development and global adoption of AI/ML-enabled technologies and secure, trusted, and open digital ecosystems that promote values critical to free and open societies;
- Elevate—across U.S. departments and agencies—the prioritization of digital development necessary to advance U.S. interests and IDDI goals and reorient U.S. development efforts for a digital age; and
- Strengthen U.S. foreign policy through significant appropriations for digital development, increased resourcing and staffing, and expanded authorities for federal departments and agencies, particularly the Department of State, USAID, and DFC.

Actions for the ISTS Task Force:

- **Develop, as part of the ISTS, a U.S. national strategy for promoting digital technologies and supporting digital development, infrastructure, and capacity-building.**
 - o The ISTS should include a comprehensive and integrated approach to the foreign assistance and development financing tools of the U.S. government. This will enable coordinated U.S. participation in the broader IDDI effort and provide a roadmap to more effectively using U.S. government resources to support digital infrastructure development and democratic adoption of AI and emerging technology.
 - o The strategy should also detail a strategic messaging and public awareness campaign to expose violations of international standards and democratic norms by authoritarian states.
 - o The figure below identifies critical U.S. stakeholders and their proposed role in the government-wide effort. The Commission recommends development of agency-specific plans to implement the strategy.

U.S. National Plan to Support International Technology Efforts

U.S. National Plan to Support International Technology Efforts.



International Science & Technology Strategy (ISTS) Task Force

Convened by the White House with leadership from State, Treasury, Commerce, Energy, DFC, EXIM, MCC, NSF, USAID, USTDA, and other critical agencies

Develop and oversee implementation of U.S. government-wide strategy for international technical standards and international digital development efforts

Select U.S. Stakeholders and Proposed Roles

 <p>Department of State <i>Foreign policy leadership and diplomacy</i></p> <p>Senior leadership and Ambassador-at-Large lead efforts to establish and implement ETC and IDDI</p> <p>Implement holistic effort to coordinate international security, economic policy, S&T, human rights, foreign assistance</p> <p>Facilitate tech diplomacy through U.S. Embassies and Missions</p>	 <p>United States Agency for International Development (USAID) <i>Digital development and humanitarian assistance</i></p> <p>Prioritize implementation of digital development through Digital Strategy</p> <p>Provide resources, tools, and expertise for broader U.S. digital development projects</p> <p>Advise on international technical standardization</p>
 <p>Department of Commerce</p>	 <p>U.S. International Development Finance Corporation (DFC) <i>Foreign direct investment</i></p> <p>Expand investments in technology and digital infrastructure</p> <p>Increase blended finance transactions to achieve scale</p>
 <p>National Institute of Standards and Technology (NIST) <i>Advise on international technical standardization</i></p> <p>Coordinate interagency task force on technical standards</p> <p>Improve partnership and collaboration with industry</p> <p>Align standards for secure, reliable, and trusted technologies with key allies and partners</p>	 <p>U.S. Export-Import Bank (EXIM) <i>Export promotion and financing assistance</i></p> <p>Leverage Program on China and Transformational Exports to strengthen U.S. tech competitiveness</p> <p>Advise on incentivizing and exporting democratic emerging technologies</p>
 <p>Bureau of Industry and Security (BIS) <i>End-user controls</i></p> <p>Develop and coordinate end-user licensing policies and export controls to further democratic values</p>	 <p>U.S. Trade and Development Agency (USTDA) <i>Export promotion and technical assistance</i></p> <p>Support export of U.S. emerging technologies through increased funding, training, assistance, and pilot projects</p>

NIST is non-regulatory agency of the Department of Commerce and BIS is a Bureau in the Department of Commerce.

- **Conduct an assessment of existing programs across the U.S. government and associated funding, staffing, and authorities of ISTS Task Force entities.**
 - o The ISTS Task Force should conduct an early assessment to guide development of this portion of the ISTS. The assessment should include:
 - An evaluation of current and recent interagency programs to identify best practices and priority countries as well as data governance frameworks and multilateral engagements on which more comprehensive efforts can be built.³¹ *See Table 3. U.S. Digital Development Programs (located at the end of this Plan).*
 - An evaluation of authorities, appropriations, and personnel necessary to achieve the objectives of the national strategy.³²
 - o The ISTS Task Force should also consider addressing immediate needs for experts in emerging technology issues through innovative public-private fellowship rotation programs and intra-government details.
 - Personnel from CISA and the Defense Advanced Research Projects Agency (DARPA), for example, can help on immediate needs at the Department of State and USAID.

Action for USAID:

- **Prioritize implementation of the Digital Strategy and support urgent resourcing and organizational needs.**
 - o The USAID Digital Strategy³³ is an ambitious and necessary five-year plan for development and humanitarian assistance focused on promoting secure, trusted, and open digital ecosystems and the responsible use of AI technologies.
 - Implementation has lagged due to insufficient funding, inadequate staffing, and bureaucratic challenges. Currently, the Digital Strategy is administered by the Technology Division within the Innovation, Technology and Research Hub in the Bureau for Democracy, Development and Innovation (DDI).
 - o The Commission recommends that the USAID Administrator continue efforts to transform the development paradigm by infusing a digital foundation across USAID portfolios.³⁴ To this end, the Administrator should prioritize the Digital Strategy by (1) advocating for congressional appropriations to fund Digital Strategy programs (see *infra*), (2) augmenting development staff with experts in AI, 5G and connectivity, and cybersecurity, both at headquarters and in forward-deployed missions, (3) converting the Technology Division into a formal Center within DDI, and (4) prioritizing the inclusion of technology and digital across all development efforts.
 - o Immediate staff augmentation can be accomplished by enhancing existing implementing partnerships with the private and non-profit sectors, through direct hires, fellowship programs for researchers, and details from other federal agencies.
 - o Longer-term staffing needs would benefit from creating a foreign service backstop from the recommended Center focused on digital expertise to strengthen USAID's ability to identify needs, assess risks, and execute on programmatic activities around digital development.

Action for DFC:

- **Expand formal relationships with international partners and private foundations to expand the scope of DFC investments and connectivity projects through blended financing arrangements.**
 - o DFC's Roadmap for Impact is a five-year effort to catalyze \$75 billion—\$25 billion by DFC, \$50 billion from the private sector—and provide technical expertise and support to optimize development impact.³⁵ The Roadmap for Impact proposes to “elevate innovation and technology across at least 50% of the DFC portfolio” and devote \$5 billion for digital infrastructure projects and increasing internet access.³⁶
 - o Current authorities limit DFC's ability to invest in higher-risk transactions, which presents challenges for scaling digital infrastructure projects, particularly in developing countries.
 - DFC investments are scored under the Federal Credit Reform Act, and DFC has limited budget authority for subsidy for equity financing (\$150 million) and debt financing and technical assistance (\$30 million).³⁷
 - DFC cannot provide concessionary lending, unlike China³⁸ and peer agencies, such as the Japan Bank for International Cooperation and European Investment Bank, as well as the World Bank.³⁹
 - o DFC should deepen its relationships with existing and new international partners to expand the scope of its financing and equity investments in the digital development space.⁴⁰
 - o Similarly, DFC should expand partnerships with a broader range of non-governmental entities to leverage its own appropriations through blended financing arrangements that enable higher-risk investments.⁴¹
 - o This may include creating a digital technology fund⁴² that invests in developing secure, trusted digital infrastructure, AI/ML-enabled technologies, and ICT with technical features that comport with democratic values and ethical norms around openness, privacy, security, and reliability.

Actions for Congress:

- **Create an allocated Emerging Technology Fund for foreign operations and related programs of USAID and the Department of State.**
 - o The underfunding of U.S. digital foreign assistance and financing programs is exacerbated by competition with other funding priorities and lack of a flexible allocated budget.
 - o Congress should authorize an allocated budget account, the Emerging Technology Fund, to facilitate holistic planning of digital foreign assistance, digital development projects, emerging technology programs, and other ISTS activities.
 - o The Commission proposes that the allocated account include the requests for additional, targeted appropriations for USAID and the Department of State.
 - o Existing digital-related programs could also be consolidated into the Emerging Technology Fund.

- **Appropriate \$200 million annually to implement the USAID Digital Strategy.**
 - The Commission recommends Congress appropriate a minimum of \$200 million annually to support implementation of the USAID Digital Strategy by the Technology Division within DDI, with required funding likely multiples higher. The funds should support programmatic activities as well as critical hiring needs.
 - This amount builds on USAID's FY 2021 request for \$82 million,⁴³ which includes support for the Digital Ecosystem Fund,⁴⁴ staff augmentation, and programmatic activities.
- **Appropriate \$300 million annually for the Department of State's emerging-technology programs.**
 - The Commission recommends Congress appropriate a minimum of \$300 million annually to support the Department of State's emerging technology programs and administrative needs and to build what is currently a small cross-Department group of officials with expertise in emerging technology issues.
 - These funds should include the immediate request for supplemental appropriations, described later in this Blueprint for Action, for \$70 million to address urgent diplomatic efforts, programs, and foreign operations in AI, emerging technologies, and data.
 - Additional funding would support foreign assistance activities around emerging tech and digital infrastructure, to include planning, assessments, and provision of assistance. Funds would support targeted, digital programs in several areas, including Department of State programs involving the rule of law (INL), democracy and human rights (DRL), security cooperation (AVC, PM, ISN), and technical assistance (EB, STAS, others).
- **Provide DFC with sufficient appropriations to strengthen development finance as a tool for achieving national objectives.**
 - To improve the ability of the U.S. government to leverage the tools of development finance and equity investments to further the ISTS mission, Congress should provide DFC with \$1 billion in flexible, programmatic funding to support digital development projects.⁴⁵
- **Increase DFC's capacity for blended development financing through interagency partnerships.**
 - Congress has restricted the appropriations available for USAID, MCC, and the Department of State to partner with DFC in blended transactions.⁴⁶ USAID and the Department of State are limited to transferring \$50 million overall—spread across all projects, not limited to digital.⁴⁷
 - As DFC's role in digital development investments increases, the need for funds from the Department of State, USAID, and MCC will also increase, requiring an equivalent increase in funding to support USAID, State, and MCC digital and AI-related efforts that may be tabled to enable a transfer of funds to DFC.⁴⁸
 - The Commission proposes that Congress appropriate a total of \$200 million to the Department of State, USAID, and MCC to be used for DFC investment programs.
- **Appropriate funds to support critical personnel needs at DFC.**

- o Congress should appropriate funds sufficient for DFC to increase its forward-deployed personnel, located in regions in which DFC invests.
 - Currently, 98% of DFC staff is based in Washington, D.C. This puts DFC at a disadvantage vis-à-vis foreign development finance institutions (DFIs). By comparison, DFC estimates that peer DFIs have roughly four times the number of staff and base them predominantly in low- and lower-middle-income countries.⁴⁹

Component 3: Promote Transparency and Accountability Through Export Controls

ISTS objectives will be furthered by the U.S. government's ability to harness the power of the U.S. private sector. A critical tool for achieving this involves incentivizing the export of technologies that align with democratic values.

Action for the Departments of Commerce and State:

- **Develop end-user licensing policies and export controls as part of the ISTS.**

- o The Department of Commerce, through the BIS, should use targeted end-use controls and human rights due-diligence reporting requirements to prevent and deter U.S. firms from enabling problematic government end uses of AI and associated technologies.⁵⁰
 - BIS should build on its 2020 request for public comments on ways to strengthen controls and monitoring of advanced surveillance systems—this area could be explored to prevent the use of compute-intensive technologies for human rights abuses while furthering the promotion of democratic-aligned technology. Regulations issued in October 2020 provide BIS with discretion to deny export licenses for products that could be used to violate or abuse human rights.⁵¹
 - Coordinated with the ISTS Task Force, these stronger export control rules can promote the ethical and responsible use of AI among U.S. firms, set standards for global industry, and counter abuses of human and civil rights.
- o The Department of State, through the Bureaus of Democracy, Human Rights and Labor (DRL), International Security and Nonproliferation (ISN), and Political-Military Affairs (PM), should expand upon its recently issued framework to guide businesses in assessing risks of human rights abuses when exporting surveillance equipment,⁵² while bolstering the promotion of democratic values.
- o In coordination with the Department of Commerce, the Department of State should expand data collection and analysis of human rights abuses associated with emerging technologies and authoritarian digital practices.⁵³

Recommendation

Recommendation: Enhance the United States' Position as an International Emerging Technology Research Hub


The third component of the ISTS is to enhance the role of the United States as an international emerging technology research hub. The goals are to:

- Facilitate U.S. government contributions to collaborative initiatives and technical standards, such as Global Partnership on AI (GPAI)⁵⁴ and digital projects of the OECD;
- Strengthen the talent of the United States, allies, and partners by investing in people through workforce development, mentorship, and exchange programs facilitated through the recommended Multilateral AI Research Institute (MAIRI);
- Foster collaborative research relationships and pool research resources for the development of technologies (particularly in civilian applications) that comport with democratic values and address gaps in commercial R&D, including joint research in privacy-enhancing technologies; and
- Enable the U.S. and allies to overcome current regulatory challenges currently inhibiting collaboration, particularly in Europe, such as data-sharing restrictions and liability agreements.

Components of International Digital Research HUB.

Support International R&D

Formalize relationship between U.S. National AI Research Institutes and the Global Partnership on AI
Support OECD-led efforts to further AI-related R&D



Establish Multilateral AI Research Institute (MAIRI)

NSF-led collaboration with allies and partners; physically located in the U.S.
MAIRI will enable equitable, multilateral research

Catalyze International Collaboration

Leverage existing visa programs to facilitate foreign researchers' travel

Component #1: Support International Digital and AI R&D

International efforts, like the GPAI and the OECD's AI and digital initiatives, are critical forums for facilitating alignment among like-minded countries on advancing the responsible and human-centric development and use of AI. Research undertaken by the National AI Research Institutes—run by the NSF and other U.S. agencies—and by other programs across Federal departments and agencies, is an incredible resource that should support these key international efforts and advance AI and digital goals of the U.S. and like-minded partners.

Actions for the Department of State, OSTP, and NSF:

- **Formalize a “center of expertise” relationship with the Global Partnership on Artificial Intelligence (GPAI).**

- o NSF should evaluate candidates to serve as a U.S. center of expertise for GPAI. NSF should submit a recommendation to the Director of OSTP to guide negotiations with GPAI.
 - NSF should consider candidates from among its AI-related awardees, including the National AI Research Institutes or by establishing a coordination hub of all Institutes. NSF should also propose methods for leveraging other U.S. science and research agencies, such as the Department of Energy and NIST, to support the center of expertise.
- o In coordination with the Department of State and OSTP, NSF should negotiate a memorandum of understanding between NSF and GPAI to formalize the center’s support of GPAI working groups.⁵⁵

- **Increase support to the OECD’s AI and digital efforts.**

- o The U.S. government should expand its collaboration with the OECD’s AI initiatives, including those of the Directorate for Science, Technology and Innovation and the AI Policy Observatory.⁵⁶ The Commission proposes an expanded relationship in three ways:
 - o First, NSF should explore methods to support OECD’s “Going Digital” program⁵⁷ to promote data sharing among partner nations.
 - Input should include pilots on cross-border data flow measurement, taxonomies to compare countries’ data initiatives, or data governance policies.
 - o Second, the Department of State, OSTP, and NSF should look for opportunities to align with allies and partners through the OECD on data guidelines, in particular by promoting value-based best practices for collecting (e.g., with consent and contributor controls), documenting (e.g., to support responsible use and quality), using data in R&D (e.g., with transparency), and then making data used in published research available to the broader research community (e.g., for reproducibility).
 - o Third, OSTP should work with the OECD to formalize a “network of research nodes” to coordinate AI and digital-related efforts and R&D centers worldwide. Policymakers and researchers would greatly benefit from a global information platform that enables easier understanding of the various AI and related initiatives and ongoing research efforts.

Action for Congress:

- **Provide administrative funding to support U.S. research contributions to GPAI.**

- o The centers of expertise that support GPAI also provide administrative and secretariat-like assistance (e.g., planning of GPAI plenaries). Congress should therefore provide additional resourcing to NSF to support the center’s development, administrative staff, and resourcing to leverage research from NSF’s AI portfolio, including the National AI Research Institutes, and from other U.S. departments and agencies as needed.

- o The Commission recommends a minimum of \$3 million over a three-year period.

Component #2: Establish the Multilateral AI Research Institute (MAIRI)

The Multilateral AI Research Institute (MAIRI) will provide a model for equitable, multilateral research, facilitate AI R&D that builds on like-minded countries' strengths, and develop the next-generation global AI workforce. With a physical center located in the United States with a virtual presence, MAIRI will enable collaborative research among key allies and partners and contribute to a broader effort—reflected in the Emerging Technology Coalition and IDDI—to preserve free and open societies, win the global technology competition, and foster AI innovation in a manner that comports with democratic values. Ultimately, to further these objectives, MAIRI should seek to facilitate a federated network of research institutes across the globe and with national labs and university hubs.

Actions for the NSF:

- **Establish MAIRI in the U.S. and support involvement of U.S. researchers in MAIRI.**
 - o NSF should establish MAIRI, modeled on the Banff International Research Station.⁵⁸ MAIRI should have a physical center in the United States, as well as a virtual presence. NSF should provide MAIRI with all staff necessary to ensure its success.
 - Although NSF does not require further authorities to establish MAIRI, legislation could facilitate this process (see actions for Congress).
 - o MAIRI should be designed with sufficient flexibility to enable involvement by researchers from industry, academia, and research institutions and philanthropies on a project-by-project basis; other U.S. departments and agencies, like the Department of Energy, may be critical for leveraging the entire U.S. R&D ecosystem.
 - o NSF programs through the Office of International Science and Engineering (OISE) can support MAIRI by facilitating involvement of U.S. researchers.
 - AccelNet⁵⁹ can fund the travel, virtual networking and other activities necessary to support research projects between research networks.
 - MULTIPLIER⁶⁰ may support subject-matter experts' travel to identify collaboration opportunities with founding members or with countries that are considering joining MAIRI.
- **Identify key allies and partners to be MAIRI founding members.**
 - o NSF, in close coordination with the Department of State, should identify and negotiate involvement of founding members.
 - o The Commission recommends that founding members include Australia, Canada, France, Germany, Italy, Japan, New Zealand, South Korea, and the United Kingdom.
 - These countries have existing agreements and collaborative relationships with the United States that could be more readily leveraged to develop the center. They also have extensive research capabilities and share values and interests with the United States.

- o Expansion to include additional allies and partners should be prioritized; for example:
 - European Union involvement should be a priority; however, the EU's inclusion in MAIRI will depend on the ability to overcome disagreements between the EU and United States over governing law, liability, funding, data sharing, and intellectual property.
 - Involvement by India should also be prioritized as MAIRI develops, building on the Commission's recommended U.S.–India Strategic Tech Alliance.

- **Develop research integrity principles with MAIRI's founding members.**

- o Founding members would agree to a jointly determined Principles for Multilateral AI Research, which would be founded on the importance of research integrity. Principles may include the need for transparency, particularly in disclosing funding and international connections; the necessity for open data and data sharing; the development of risk–benefit frameworks; and the use of merit-based competition reviews of research proposals.
- o Members would also receive training on security risks and agree to use trusted infrastructure as part of founding principles (see recommended appropriations in actions for Congress).
- o The agreement will also detail the terms for handling intellectual property, sharing data, governing law and liability, and funding.

- **Develop a concrete research agenda with MAIRI's founding members.**

- o Once founding members have agreed to the Principles, they will determine focus areas and initiatives. Countries will fund the involvement of their researchers in joint projects. Joint research projects will occur through virtual spaces as well as at partner entities like research institutions and universities that receive funding from MAIRI. The facilities of other participating departments and agencies may also be used.⁶¹
- o Research Priorities: Projects should be chosen to leverage members' comparative advantages, enabling participants to learn from partner researchers. Examples of R&D priorities are provided in the Emerging Technology Coalition Annex to this Blueprint for Action. Priorities should include:
 - Building shared, secure compute resources (including high–performance computing [HPC], cloud, and quantum computing),⁶² including joint benchmarking projects and data–sharing, pooling, and storing initiatives founded on commonly agreed upon principles that ensure trust, privacy, and security.
 - Privacy–preserving AI/ML technologies, including technologies like federated learning and on–device prediction that enable remote execution, encrypted computation through multi–party computation and homomorphic encryption, and differential privacy.
 - Developing smart–city technologies, aligned with democratic values, that promote sustainability as well as norms that should guide standards development at bodies like the ITU and technical standards bodies.

- **Coordinate with MAIRI founding members on funding, international agreements, and governance structures.**
 - o Although the United States should fund the initial startup costs, including acquisition of MAIRI's physical center, staff, and virtual research/networking infrastructure (see recommendations to Congress below), each member should thereafter provide proportionate financial contributions to MAIRI's R&D and to the participation of their researchers in MAIRI-sponsored workshops and conferences⁶³ modeled on the approach used by the Banff International Research Station.⁶⁴
 - o For ongoing operations, MAIRI should explore the potential to develop an endowment, modelled on the three US-Israeli binational funds. This approach would facilitate the use of philanthropic donations to support MAIRI.
 - o Umbrella international AI/S&T agreements—negotiated with NSF, MAIRI members, and U.S. agencies—will facilitate cooperation among allies and partners beyond MAIRI.
 - o Once established, MAIRI may support GPAI and other international efforts. MAIRI should also pursue research agreements with other centers of excellence and research centers focused on AI R&D to create a federated network of research institutes throughout the globe.
 - o MAIRI members should also determine how they will determine expanding MAIRI's membership, particularly to the European Union and India.

Action for Other U.S. Departments and Agencies:

- **Support the establishment of MAIRI and its R&D.**
 - o NSF will be the U.S. anchor partner for MAIRI. Its success requires leveraging the entire U.S. R&D ecosystem and government research entities.⁶⁵ The Departments of Energy and State as well as NIST, in particular, should be critical partners.
 - o The Department of Energy should leverage its national labs, history of working with industry, immense technical capabilities, experience on applied research, and expertise in HPC and quantum computing.
 - o The Department of State should provide foreign policy expertise and diplomacy, including by assigning a dedicated Foreign Service Officer to support the creation of MAIRI as well as identification of beneficial projects.
 - o Other federal entities, including the National Institutes of Health, the National Oceanic and Atmospheric Administration (NOAA), and the Department of Health and Human Services, will be critical for technical expertise and collaboration on targeted research projects.

Actions for Congress:

- **Pass legislation to formally authorize MAIRI.**
 - o Although it's not required for MAIRI's establishment, Congress should pass legislation that formally authorizes the creation of MAIRI and clarifies the authorities

of other executive agencies to award funding to MAIRI. This will serve as a signal of the importance of international AI collaboration and ensure that NSF and partner agencies have sufficiently robust authorities to achieve its objectives.

- o Legislation should also specifically authorize and direct NSF, in coordination with the Department of State, to create a trusted learning cloud and associated compute capacity to facilitate international collaborative research.
 - The trusted learning cloud would enable access to needed resources, compute, and data for shared innovation and development of data-sharing standards that could be a model for a larger international data-sharing framework.
- **Support the establishment of MAIRI through appropriate funding to NSF and other critical agencies.**
 - o The Commission recommends Congress appropriate a minimum of \$60,750,000 for a five-year period, which will be supplemented by contributions from international partners.
 - o The proposed appropriations are as follows:
 - \$10 million per year for five years to NSF and other critical agencies (such as the Departments of Energy and State) for research initiatives.
 - \$2 million per year for five years to NSF for establishing and maintaining the physical center located in the United States, its associated infrastructure, and administrative operations.
 - \$150,000 per year for five years to NSF to support U.S. researchers' travel and associated expenses to partake in MAIRI's workshops, conferences, and other events at the physical center.
 - o The Commission recommends Congress appropriate \$11.25 million per year for research initiatives dedicated to creating a trusted learning cloud and associated compute capacity to facilitate international collaborative research.
- **Create an endowment for MAIRI to support ongoing funding.**
 - o MAIRI may wish to develop an endowment fund similar to the U.S.-Israeli binational foundations. If pursued, Congress should authorize this endowment fund and support an initial U.S. investment. Additional appropriations would be required to support a MAIRI endowment fund secretariat.

Component #3: Expand Talent Exchanges

The United States must attract talent to collaborative research endeavors at both the National AI Research Institutes and MAIRI. Sustained, strong collaboration between MAIRI partners is critical to ensure that responsible, secure, human-centric AI prevails over authoritarian AI. Shoulder-to-shoulder research and talent exchanges are invaluable, enabling researchers to build relationships, learn from each other, exchange ideas, and spark future collaborations.

Action for the Department of State:

- **Leverage O and J visa programs to attract skilled researchers to support MAIRI and international talent exchange programs.**
 - o The Department of State, in coordination with the Department of Homeland Security, should leverage the O and J visa programs to facilitate foreign researchers to travel to the United States to work collaboratively with researchers from the United States and other nations.⁶⁶ There are no statutory caps on the number of visas issued under these programs.⁶⁷

Recommendation

Recommendation: Reorient U.S. Foreign Policy and the Department of State for Great Power Competition in the Digital Age

In the near term, it is imperative to establish a Department of State focal point for emerging technology policy and expertise and resourcing through steps the Commission proposes below. In the longer term, the United States must fundamentally reorganize the structure, focus, and culture of the Department of State to advance American interests at the intersection of democracy, technology, security, commerce, and human rights.⁶⁸ Without high-level support in the Department, technology competition is unlikely to become a core aspect of U.S. foreign policy.

Action for the President:

- **Disseminate a Presidential letter of instruction to Chiefs of Mission that articulates emerging technology as inseparable from U.S. core geopolitical interests.**
 - o The instruction should direct each Chief of Mission to develop an emerging technology plan as part of its mission strategy submitted to the Secretary of State.

Actions for the Department of State:

- **The Secretary of State should direct the Deputy Secretary of State for Management and Resources (D/MR) to lead on reorienting and reorganizing the Department for technology diplomacy.**
 - o The D/MR position has in the past exercised leadership to oversee significant organizational and resourcing priorities across the Department of State.
 - Past officials in the D/MR position have spearheaded U.S. diplomatic priorities around regional policy (such as the U.S.–Pakistan Strategic Dialogue), foreign assistance, civilian response, and international economic issues.
 - o D/MR should provide direction around immediate and long-term planning to coordinate disparate offices and bureaus within the Department of State, develop technological expertise at all levels of the Foreign and Civil Service, and ensure that policy direction is aligned with management, personnel, and resource actions needed to achieve reorientation with urgency and sustainability.
 - o D/MR should also provide leadership for executing the ISTS.

- **Generate a comprehensive proposal for immediate funding needs with a request to Congress for supplemental appropriations.**
 - The Department of State should prepare and submit to Congress within 60 days a request for immediate funding needs to address personnel shortages and programmatic efforts to further U.S. diplomacy around emerging technology. The Department should seek funding through supplemental appropriations to avoid lags in the budget cycle.
- **Expedite building out a dedicated bureau for emerging technology diplomacy.**
 - The Department of State should expedite and prioritize efforts to staff, resource, and build out a bureau for emerging technology diplomacy.
 - The Bureau of Cyberspace Security and Emerging Technologies (CSET Bureau), formally approved in January 2021,⁶⁹ is intended to focus on security challenges associated with cyberspace and emerging technologies.⁷⁰ The Commission proposes that the CSET Bureau be established with a broad aperture to address diplomatic efforts across the security, economic, human rights, and regional dimensions of foreign policy. It should serve as a clearinghouse to assess strategic, budgetary, and personnel priorities on emerging technology policy across the Department. The Bureau should have responsibilities for managing high-level dialogues with allies and partners to further progress and cooperation, coordinating policy, standards, and digital development assistance with U.S. agencies, and promoting AI and emerging technology advocacy within the Department.
 - The Department should assess where the CSET Bureau should be placed to best achieve those objectives, but must ensure that its creation is not further delayed.
 - The Bureau should be led by a high-profile Assistant Secretary or Ambassador-at-Large. If the Department appoints an Assistant Secretary to head the Bureau and lead coordination across the Department, it should consider creating a separate Ambassador-at-Large position to lead diplomacy with foreign counterparts on cybersecurity and emerging technology.
- **Develop a comprehensive plan to reorganize technology diplomacy under a new Under Secretary.**
 - The Department of State should develop a comprehensive proposal to establish an Under Secretary for Science, Research and Technology (Q). State/Q would bring together the elements for a robust, coordinated approach to science and technology diplomacy in the context of great power competition—with a focus on emerging technology.⁷¹
 - State/Q would also work with the Director for Foreign Assistance to manage a new allocated account for digital democracy and emerging technologies and lead implementation of the ISTS across the U.S. government.
 - The plan should also consider establishing Deputy Assistant Secretaries for Science and Technology in each regional bureau. These positions would provide a critical link between technology officers and senior leadership.
 - Currently the Department lacks a core of senior, career officials with deep and broad technology policy expertise. The positions would provide a career path to the senior level for officers focused on technology policy and would enable senior-level advocacy for reforms needed to effectively manage technology policy.

- o Given the urgent need to enhance the Department's technology diplomacy capacity and the likely long-term nature of the process of establishing a new Under Secretary, efforts to implement this recommendation should proceed in parallel with the Commission's other organizational recommendations.
- **Establish a diplomatic presence in major U.S. and foreign technology hubs.**
 - o The Department of State should enhance its presence in major foreign and U.S. technology hubs, supported by establishing a cadre of dedicated technology officers at U.S. missions to strengthen diplomatic advocacy, improve technology scouting, and inform policy and foreign assistance choices.
 - o The Department should accelerate plans to establish 12 Regional Technology Officer positions around the world⁷² and further describe how these officers will enhance U.S. technology competitiveness with partners such as the Foreign Commercial Service, USAID, and DFC. These officers should also scout technology initiatives that can enhance our diplomatic and development capabilities.
 - o The Department should re-establish a permanent presence in Silicon Valley, which it initiated in 2014, and established dedicated positions in 2015–2016. These positions and State's presence were discontinued when an OMB hiring freeze was implemented in January 2017.
 - o In addition, the Diplomat in Residence program—with presence in 16 regions at universities across America⁷³—should be repurposed beyond recruitment to include public diplomacy, technology scouting, and engagement with foreign government and commercial entities active across America. Domestic insight is a valuable input into foreign policy and will increase public confidence in and support for America's international technology leadership.
- **Incorporate AI and emerging technology training modules into Foreign Service institute (FSI) courses.**
 - o The Department of State should incorporate mandatory AI and emerging technology—related modules into key FSI training courses, including the Ambassadorial Seminar, the Deputy Chiefs of Mission course, Political and Economic Tradecraft courses, and A-100 orientation training classes. FSI should also develop a stand-alone course on emerging technologies and foreign policy.
 - o The Department should partner with academic and private-sector organizations to access the leading edge of technology education while also building a more robust technology fellows program for exchanges with industry.

Actions for Congress:

- **Expedite necessary reorganization of the Department of State by passing legislation to create an Under Secretary for Science, Research and Technology (Q).**
 - o Congress should act to create the State/Q position and consolidate disparate S&T efforts in the Department in a single division. There is urgent need for such reorganization, and Congress can empower the Department of State by introducing and passing legislation to expedite the reorientation.
- **Appropriate funds to support immediate augmentation of the U.S. diplomatic corps.**

- o Congress should provide robust funding for hiring and training of needed personnel to enable the Department of State's reorientation and support the Department's international efforts to promote U.S. values and standards in AI, data, and associated emerging technologies.
- o The Commission recommends that Congress provide at least \$8 million in supplemental appropriations for immediate hiring of staff to address emerging technology needs across the Department's offices and bureaus, to establish a diplomatic presence in major U.S. and foreign technology hubs, and to develop FSI courses.
- o See the Funding Table Appendix to this report for a detailed breakdown of the proposed appropriations.

- **Appropriate funds to support the CSET Bureau.**

- o The Commission recommends a minimum of \$20 million to establish the CSET Bureau.
- o See the Funding Table Appendix to this report for a detailed breakdown of the proposed appropriations.

- **Appropriate funds to support critical diplomatic efforts, programs, and foreign operations in AI, emerging technologies, and data.**

- o Further funding is needed to enable the Department of State to advance responsible AI aligned with U.S. and like-minded values.
- o While details of funding needs should reflect input from the Department of State, the Commission recommends, at a minimum, that Congress issue a supplemental appropriation for no less than \$37 million, as a subset of the proposed \$300 million described earlier in this Blueprint for Action, for the following urgent needs:
 - Public diplomacy messaging and engagement to support democratic values and raise awareness of U.S. leadership in AI innovation as well as the risks of unwanted technology transfer and authoritarian digital practices;
 - AI exchange programs to promote U.S. values and fund participation by developing countries participation in multilateral AI activities;
 - Programs to showcase American innovation and through which to promote the ethical use of AI, including the American Spaces, TechCamp, MakerSpaces, and U.S. Speakers programs⁷⁴;
 - Partnership development and advancement of scientific norms through the U.S. Science Envoys and Embassy Science Fellows programs;
 - Diplomatic operations and programs around international AI cooperation, including support for initiatives of the ETC;
 - Promotion of human rights and fundamental freedoms in the AI context;
 - Maintaining U.S. lead in the use of AI for military applications through cooperation with allies and partners;
 - Ensure political and policy congruence with allies and partners on the use of AI-enabled weapon systems;

- Ensure continued interoperability among the U.S. and its allies and partners;
- Training and capacity-building for foreign governments on emerging technologies to support responsible innovation;
- Reporting to counter malign influence in AI ecosystems;
- Empower global AI-focused S&T entrepreneurship through the U.S. Global Innovation through Science and Technology (GIST) Initiative⁷⁵; and
- Public diplomacy initiatives on international AI standards as well as tracking and reporting on public opinion related to AI.
- See the Funding Table Appendix to this report for a detailed breakdown of the proposed appropriations.

Table 1. Key Multilateral Technology Initiatives

	Initiatives	Critical Areas	Objectives
Multilateral AI and Emerging Technology Initiatives	Council of Europe/CAHAI	<ul style="list-style-type: none"> • Standards and Norms • Promote Human Rights and Democracy 	<ul style="list-style-type: none"> • Support the Ad Hoc Committee on Artificial Intelligence (CAHAI) efforts around the development and use of AI aligned with human rights, rule of law, and democracy (goal of international legislative framework on AI similar to the Budapest Convention on Cybercrime) • Membership: 47 countries from Europe and six Observer States (Canada, US, Israel, Japan, Mexico, Holy See)
	D10 Initiative	<ul style="list-style-type: none"> • Standards and Norms • Promote Human Rights and Democracy 	<ul style="list-style-type: none"> • Foster international cooperation to provide 5G alternatives to ZTE and Huawei; shift critical supply chains out of China, and protect national security • Support nascent effort as it builds on the promising coalition and refines its goals, structure, and timeline; explore potential for the UK and U.S. to jointly announce further developments • Explore expansion into other emerging technologies critical to national security • Additional consideration: Divergent views on 5G and absence of key nations in 5G effort may limit efficacy • Membership: Australia, Canada, France, Germany, India, Italy, Japan, South Korea, the UK (founding member), and the U.S.
	OECD Digital and AI efforts (OECD AI Principles, OECD.AI, ONE AI, “Going Digital” Project)	<ul style="list-style-type: none"> • Standards and Norms • Data Sharing • IDDI • Joint R&D on AI and Digital Infrastructure 	<ul style="list-style-type: none"> • Continue articulating support of OECD’s international efforts to advance responsible AI and promote implementation of OECD AI Principles • Advance shared interests through AI Policy Observatory (OECD.AI) and the Network of Experts on AI (ONE AI), particularly in working groups on classifying AI systems, implementing values-based principles, and guiding national AI strategies • Facilitate coordination of U.S. experts and representatives engaging in ONE AI and associated working groups • Advance U.S. interests in “Going Digital” initiatives that promote data sharing and harmonizing on IP and regulation • Membership: 37 member countries, OECD AI Principles adopted by 42 countries

	Initiatives	Critical Areas	Objectives
Multilateral AI and Emerging Technology Initiatives	Freedom Online Coalition (FOC)/ T-FAIR	<ul style="list-style-type: none"> Standards and Norms IDDI 	<ul style="list-style-type: none"> Engage to further efforts on AI and human rights and against digital authoritarianism, with a focus on content moderation and facial recognition Support Task Force on AI and Human Rights (T-FAIR) Membership: 32 country coalition with members from Africa to Asia, Europe, the Americas, and the Middle East
	GPAI	<ul style="list-style-type: none"> Standards and Norms Data Sharing Joint R&D on AI and Digital Infrastructure IDDI 	<ul style="list-style-type: none"> Influence direction, scope, and goals; support expanded membership Formalize a U.S.-based center of expertise to provide technical support for working groups Advance shared interests of democratic nations through working groups: 1) Responsible AI (including Ad Hoc AI and Pandemic Response Subgroup), 2) Data governance, 3) Future of work, 4) Commercialization and Innovation and facilitate coordination of U.S. experts and representatives engaging in working groups and steering committee Engage in Multi-stakeholder Experts Group Plenary Membership: Australia, Canada, France, Germany, India, Italy, Japan, Mexico, New Zealand, Singapore, Slovenia, South Korea, Spain, the United Kingdom, the United States, and the European Union with the OECD and UNESCO as a Permanent Observer the GPAI Council and Steering Committee
	IPAC	<ul style="list-style-type: none"> Promote Human Rights and Democracy Standards and Norms Promote and Protect Innovation IDDI Defense and Security* 	<ul style="list-style-type: none"> Influence IPAC to become a vehicle for promoting AI-related goals of democratic nations, including the spread of democratic technology alternatives; alignment on emerging technologies that pose threats to national security; alignment on fair trade practices for digital commerce; countering IP theft; alignment on export controls and investment screening; and overall legislative alignment Membership: IPAC Co-chairs span the political spectrum and come from 19 countries across North America, Europe, Asia, and Africa
Security Alliances and Partnerships	Five Eyes/ TTCP AI Strategic Challenge	<ul style="list-style-type: none"> Defense and Security 	<ul style="list-style-type: none"> Develop methods to address AI application and interoperability, including a possible test bed for application in other situations and with other coalitions (e.g., NATO) Membership: Australia, Canada, New Zealand, the United States, and the United Kingdom
	JAIC AI Partnership for Defense	<ul style="list-style-type: none"> Defense and Security Data Sharing Standards and Norms 	<ul style="list-style-type: none"> Strengthen role as key multilateral forum for security coordination around AI, including advancing shared interests and best practices on AI ethics, collaborative frameworks, and strategic messaging Current membership: United States plus twelve partner nations—Australia, Canada, Denmark, Estonia, Finland, France, Israel, Japan, Norway, South Korea, Sweden, and the United Kingdom
	NATO	<ul style="list-style-type: none"> Defense and Security Data Sharing Standards and Norms 	<ul style="list-style-type: none"> Promote interoperability, human capital development, implementation of strategic objectives Membership: 30 countries from Europe and North America
	National Technology and Industry Base (NTIB)	<ul style="list-style-type: none"> Defense and Security Promote and Protect Innovation 	<ul style="list-style-type: none"> Leverage NTIB to strengthen the industrial capabilities of the U.S. and allies and address supply chain concerns; explore expansion of mandate to advance AI-related interests around defense and security issues Membership: Australia, Canada, the United Kingdom, and the United States

*In addition to mapping multilateral efforts on AI and associated technologies to the seven critical areas recommended as priorities for the Emerging Technology Coalition, and detailed in the Annex: Emerging Technology Coalition Blueprint, this table includes multilateral efforts that are important for defense and security, particularly as articulated in Chapter 3: AI-Enabled Warfare.

	Initiatives	Critical Areas	Objectives
Intergovernmental Forums	The Quad	<ul style="list-style-type: none"> Defense and Security 	<ul style="list-style-type: none"> Build on the Quad framework to deepen AI cooperation and negotiate formal AI cooperation agreements in the Indo-Pacific Membership: Australia, India, Japan, the United States
	Wassenaar Arrangement	<ul style="list-style-type: none"> Defense and Security Promote and Protect Innovation Promote Human Rights and Democracy 	<ul style="list-style-type: none"> Advance multilateral coordination on export controls on conventional and dual-use technologies Membership: 42 participating states—Australia, Argentina, Canada, India, Japan, Mexico, New Zealand, Russia, South Africa, South Korea, Turkey, Ukraine, the United Kingdom, the United States, and all EU members other than Cyprus
	G7	<ul style="list-style-type: none"> Standards and Norms Data Sharing Promote and Protect Innovation 	<ul style="list-style-type: none"> Support French-led effort to use G7 to promote responsible AI development and coordinate on efforts to counter disinformation and other dangerous online content Additional consideration: Success of D10 as a coalition to address common AI-related issues may limit G7 efficacy in the area, although it continues to serve as a key forum to address important geopolitical topics Membership: Canada, France, Germany, Italy, Japan, the United Kingdom, and the United States
	G20	<ul style="list-style-type: none"> Joint R&D on AI and Digital Infrastructure Standards and Norms Data Sharing 	<ul style="list-style-type: none"> Advance efforts to enable international digital economy, develop global solution to tax challenges from digitization of the economy, and utilize technology in infrastructure and smart city efforts Ensure countries do not successfully promote authoritarian technology, particularly on topic of smart cities Membership: Governments of 19 countries and the EU; includes China, Russia
	IP5	<ul style="list-style-type: none"> Promote and Protect Innovation 	<ul style="list-style-type: none"> Continue to engage in IP5's New Emerging Technologies AI Task Force to advance global legal certainty and protections of AI-related IP, enhance efficiencies in office operations through AI adoption, and strengthen communication with industry Membership: European Patent Office and national patent offices from Japan, South Korea, China, and the United States
	International Standards	3GPP	<ul style="list-style-type: none"> Standards and Norms Promote and Protect Innovation
IEEE		<ul style="list-style-type: none"> Standards and Norms Data Sharing 	<ul style="list-style-type: none"> Engage, particularly on standards within the P7000 series on ethically aligned design series (e.g., P7001 - Transparency of autonomous systems and P7003 - Algorithmic Bias) and Ethics Certification Program for Autonomous and Intelligent Systems (ECPAIS) Membership: IEEE is a technical professional organization that has 419,000 members from more than 160 countries
ISO/IEC		<ul style="list-style-type: none"> Standards and Norms Data Sharing 	<ul style="list-style-type: none"> Advance standards to enable innovation; protect national and economic security Maintain consensus approach to standards development; counter adversarial or politicization efforts Ensure U.S. domestic policy and resourcing enables full U.S. engagement, particularly in ISO/IEC JTC 1/SC 42 on AI Membership: 31 participating member countries and 16 observing member countries/territories

	Initiatives	Critical Areas	Objectives
International Organizations	ITU-T	<ul style="list-style-type: none"> Standards and Norms Data Sharing 	<ul style="list-style-type: none"> Continue to engage on international technical standards and defend the integrity of international technical standards Engage in AI for Global Good Summit, which has strong involvement from India, China, and Japan Membership: 193 countries are members of ITU-T
	UNESCO	<ul style="list-style-type: none"> Standards and Norms 	<ul style="list-style-type: none"> Engage with the Ad Hoc Expert Group's initiative to create a global standard-setting instrument on ethics and AI (Recommendation on Ethics of AI) Support continued development of the AI Decision Makers' Essential Toolkit to help decision makers, particularly in Africa, address practical questions for AI and UNESCO's AI capacity building programs for stakeholders in the judicial system Membership: 193 member countries
	UN CCW GGE on LAWs	<ul style="list-style-type: none"> Standards and Norms 	<ul style="list-style-type: none"> Advance shared interests of democratic nations regarding lethal autonomous weapon systems Membership: Over 80 countries have participated in discussions
	UNSG High-Level Panel on Digital Cooperation	<ul style="list-style-type: none"> Standards and Norms Promote Human Rights and Democracy 	<ul style="list-style-type: none"> Engage as part of UN engagement; however, it is unlikely the High-level Panel on Digital Cooperation (HLPDC) will be a main vehicle for advancing U.S. AI interests Membership: Panel has participants from over 15 countries
	WIPO AI and IP Conversations	<ul style="list-style-type: none"> Promote and Protect Innovation Data Sharing 	<ul style="list-style-type: none"> Continue to engage in WIPO's "Conversations" on AI and IP Policy and Administration; includes data protection and sharing standards Membership: 193 member countries, with participants in the conversations on IP and AI receiving participants from over 130 countries
	WTO	<ul style="list-style-type: none"> Promote and Protect Innovation Data Sharing 	<ul style="list-style-type: none"> Continue to engage in e-commerce and trade efforts outflowing from the 2019 Joint Statement on Electronic Commerce and the "Osaka Track," which promotes international rule-making to promote e-commerce and addresses data concerns Membership: Although there are 164 member countries in WTO, 78 members signed the Joint Statement on Electronic Commerce; 24 countries signed the Osaka Declaration on the Digital Economy

Table 2: International Digital Development Programs

Multilateral Initiatives	<p>Trilateral Infrastructure Partnership (TIP). Established in November 2018 by the US, Japan, and Australia, the TIP supports infrastructure projects for “an Indo-Pacific region that is free, open, inclusive, prosperous, and secure.” TIP coordinates separate financing and investment efforts of Australian Infrastructure Financing Facility for the Pacific, Japan’s Bank for International Cooperation (JBIC), Japan’s Nippon Export and Investment Insurance, and U.S. government stakeholders (DFC, USAID, Department of the Interior Compact Funding, and the U.S. government’s Transaction Advisory Fund). In October 2020, TIP facilitated a \$30 million project to construct underseas fiber optic cable to Palau to ensure secure digital connectivity.</p>
	<p>Blue Dot Network (BDN). Launched in November 2019 by the US, Japan, and Australia, BDN convenes “governments, the private sector, and civil society to promote high-quality, trusted standards for global infrastructure development in an open and inclusive framework.” It leverages “commonly accepted principles and standards to promote market-driven, transparent, and financially sustainable infrastructure development” in the Indo-Pacific region and globally. It is led by the U.S. DFC, Australia’s Department of Foreign Affairs and Trade, and the JBIC.</p>
	<p>World Bank - Digital Development Partnership (DDP). Established in 2016 to support the implementation of the UN’s Sustainable Development Goals, the DDP has over 50 active client countries with an emphasis on delivering data and indicators, digital economy enabling environment, cybersecurity, accessible internet, digital government, and mainstreaming digital services, solutions, and platforms.</p>
	<p>Three Seas Initiative Investment Fund (3SIIF). The 3SIIF is a partnership of 12 countries focused on promoting economic growth, security, and a more cohesive and stronger Europe through infrastructure development in the energy, transportation, and digital sectors. The Three Seas Initiative Investment Fund (3SIIF) (currently over \$900 million, including a \$300 million DFC initial investment) supports digital infrastructure projects to meet regional compute, storage, and connectivity needs.</p>
Initiatives Led by European Partners	<p>EU Economic and Investment Plan for the Western Balkans. Announced in October 2020, this plan supports the region’s digital and green transitions while boosting economic development, regional cooperation, and EU integration. Through investment of 9 billion euros and the Western Balkans Guarantee Facility, which seeks to mobilize 20 billion euros in public and private investment, it aims to bolster digitalization across ten investment flagships. Projects include broadband infrastructure in six Western Balkan partners, establishing trustworthy data centers and cloud infrastructures, expanding the Balkan Digital highway, and digital education programs.</p>
	<p>EU-Asia Connectivity Strategy. In 2018, the European Parliament, the European Council, the European Economic and Social Committee, the Committee of the Regions and the European Investment Bank issued a joint communication to establish building blocks on an “EU Strategy on Connecting Europe and Asia” through interoperable transport, energy and digital networks. Efforts are currently underway to strengthen EU-Asia connectivity that is sustainable, comprehensive and adherent to the international rules-based order.</p>
	<p>FAIR Forward - Artificial Intelligence for All. The German Development Cooperation initiative seeks to create a more “open, inclusive and sustainable approach to AI,” particularly in Africa and Asia. In partnership with Ghana, Rwanda, South Africa, Uganda, and India, FAIR Forward seeks to: 1) strengthen local technical AI know-how, 2) remove entry barriers to AI through access to training data and AI technologies, and 3) develop policy frameworks for ethical AI, data protection, and privacy. Projects include an African AI start-up accelerator, a Pan-African digitalization initiative, and the Open for Good Alliance to improve localized AI training data.</p>
Initiatives Led by Asian Partners	<p>International Digital Cooperation - ICT Standardization (InDiCo). Launched in 2018 by the European Commission, InDiCo is a three-year project to promote ICT standards alignment and interoperability with key partner countries, including Japan, South Korea, China, India, Brazil, and the United States. It coordinates technical standardization meetings and organizes technical and political workshops on ICT standards.</p>
	<p>Quality Infrastructure Initiative. Japan has championed efforts to promote “quality infrastructure” investments, including through financing of approximately \$200 billion. Japan’s efforts have led to the development of principles for infrastructure investments and is expanding into digital connectivity.</p>
	<p>Global Cooperation and Training Framework (GCTF). GCTF, led by Taiwan and in partnership with the United States and Japan, serves as a platform for Taiwan to share its expertise with partners around the world. Recently, GCTF held a virtual webinar with Latin American and Caribbean governments on digitization, particularly on ways to leverage data and AI to help governments respond to COVID-19.</p>

Fact Sheet, U.S. Department of State, *The United States Partners with Australia and Japan to Expand Reliable and Secure Digital Connectivity in Palau* (Oct. 29, 2020), <https://2017-2021.state.gov/the-united-states-partners-with-australia-and-japan-to-expand-reliable-and-secure-digital-connectivity-in-palau/index.html>; Fact Sheet, U.S. Department of State, *2020 Indo-Pacific Business Forum Promotes Free and Open Indo-Pacific* (Oct. 29, 2020), <https://2017-2021.state.gov/2020-indo-pacific-business-forum-promotes-free-and-open-indo-pacific/index.html>; *The Launch of Multi-Stakeholder Blue Dot Network*, U.S. International Development Finance Corporation (Nov. 4, 2019), <https://www.dfc.gov/media/opic-press-releases/launch-multi-stakeholder-blue-dot-network>; *Blue Dot Network*, U.S. Department of State (last accessed Feb. 16, 2021), <https://www.state.gov/blue-dot-network/>; *Blue Dot Network: Frequently Asked Questions*, U.S. Department of State (last accessed Feb. 16, 2021), <https://www.state.gov/blue-dot-network-frequently-asked-questions/>; *Digital Development Partnership (DDP)*, The World Bank (last accessed Feb. 16, 2021), <https://www.worldbank.org/en/programs/digital-development-partnership>; *Digital Development Partnership Annual Report: Responding to the COVID-19 Crisis*, The World Bank (Oct. 26, 2020), <https://www.worldbank.org/en/news/feature/2020/10/26/digital-development-partnership-annual-report-responding-to-the-covid-19-crisis>; *Objectives*, Three Seas Initiative (last accessed Feb. 16, 2021), <https://3seas.eu/about/objectives>; *Three Seas Story*, Three Seas Initiative (last accessed Feb. 16, 2021), <https://3seas.eu/about/threeseasstory>; *The Three Seas Fund Makes Its First Digital Investment*, Three Seas Initiative Investment Fund (Dec. 2, 2020), <https://3siif.eu/news/the-three-seas-fund-makes-its-first-digital-investment>; *The Fund*, Three Seas Initiative Investment Fund (last accessed Feb. 16, 2021), <https://3siif.eu/the-fund/>; Press Release, U.S. International Development Finance Corporation, *DFC Approves Over \$2.1 Billion in New Investments for Global Development* (Dec.10, 2020), <https://www.dfc.gov/media/press-releases/dfc-approves-over-21-billion-new-investments-global-development>; Press Release, European Commission, *Western Balkans: An Economic and Investment Plan to Support The Economic Recovery and Convergence* (Oct. 6, 2020), https://ec.europa.eu/commission/presscorner/detail/en/ip_20_1811; *Questions and Answers: Economic and Investment Plan for the Western Balkans*, European Commission (Oct. 6, 2020), https://ec.europa.eu/commission/presscorner/detail/en/qanda_20_1819; *Connecting Europe and Asia - Building Blocks for an EU Strategy*, European Commission (Sept. 19, 2018), https://eeas.europa.eu/sites/eeas/files/joint_communication_-_connecting_europe_and_asia_-_building_blocks_for_an_eu_strategy_2018-09-19.pdf; *Connecting Europe & Asia: The EU Strategy*, European External Action Service (Sept. 26, 2019), https://eeas.europa.eu/headquarters/headquarters-homepage/50699/node/50699_en; *FAIR Forward - Artificial Intelligence for All*, Deutsche Gesellschaft für Internationale Zusammenarbeit (GIZ) GmbH (Jun. 2020), https://toolkit-digitalisierung.de/app/uploads/2020/06/Factsheet-FAIR-Forward_E050620-1.pdf; *Goals*, Open for Good (last accessed Feb. 16, 2021), <https://www.openforgood.info/>; *About the Project*, The InDiCo Project (last accessed Feb. 16, 2021), <https://www.indico-ictstandards.eu/about-the-project>; *The 'Expanded Partnership for Quality Infrastructure' Initiative Directed Toward the G7 Ise-Shima Summit Meeting Announced*, Japanese Ministry of Economy, Trade and Industry (May 23, 2020), https://www.meti.go.jp/english/press/2016/0523_01.html; Tobias Harris, 'Quality Infrastructure': *Japan's Robust Challenge to China's Belt and Road*, War on the Rocks (Apr. 9, 2020), <https://warontherocks.com/2019/04/quality-infrastructure-japans-robust-challenge-to-chinas-belt-and-road/>; Andreea Brînză, *Japan's Belt and Road Balancing Act*, The Diplomat (Nov. 8, 2018), <https://thediplomat.com/2018/11/japans-belt-and-road-balancing-act/>; *Global Cooperation and Training Framework (GCTF) Programs*, American Institute in Taiwan (last accessed Dec. 28, 2020), <https://www.ait.org.tw/our-relationship/global-cooperation-and-training-framework-programs-gctf/>; Fact Sheet, U.S. Department of State, *U.S. Support for Digital Transformations in Latin America and the Caribbean* (Nov. 10, 2020), <https://2017-2021.state.gov/u-s-support-for-digital-transformation-in-latin-america-and-the-caribbean/index.html>.

Table 3: U.S. Digital Development Programs

Interagency Initiatives	<p>Infrastructure Transaction and Assistance Network (ITAN) and Transaction Advisory Fund (TAF). ITAN and TAF are multi-agency, United States Government efforts to strengthen capacity-building programs, provide expert transaction advisory services, and coordinate U.S. assistance support for sustainable, transparent, high-quality digital and non-digital infrastructure in the Indo-Pacific region.</p>
	<p>US-ASEAN Smart Cities Partnership (USASCP). Launched in 2018, USASCP aims to promote U.S. private sector engagement in smart, sustainable city solutions, sharing best practice and technical collaboration, and strengthening digital economy and cybersecurity. It draws on capabilities of the Departments of State, Commerce, and Transportation as well as USAID, USTDA, and the NSF. Limited funding (\$10 million initial investment in 2018) has limited its impact.</p>
	<p>Digital Connectivity and Cybersecurity Partnership (DCCP). DCCP draws on tools of 10 U.S. agencies to promote expanded connectivity and an open, interoperable, secure, and reliable internet. In the Indo-Pacific, \$50 million has supported 15 countries through capacity building to improve digital connectivity; strengthen the private sector’s digital capacity; and assist in the design, development, and implementation of ICT policies and regulations. In the Western Hemisphere, \$10 million will support similar initiatives.</p>
	<p>Bilateral Memorandums of Understanding and Investments. United States Government entities are undertaking various bilateral projects to advance a democratic digital ecosystem. For example, EXIM and the Brazilian Ministry of Economy signed a MOU to deploy up to \$1 billion in financing to support U.S. exports to Brazil in the 5G space. DFC is financing a major telecommunications project in Ecuador and Peru, which will deploy at least 500 telecom towers and expand access to 4G mobile broadband and high-speed internet in rural areas.</p>
Department or Agency-Specific Initiatives	<p>USAID’s Digital Strategy. USAID has a far-reaching plan for digital capacity building and foreign assistance, especially around strengthening open, secure, and inclusive digital ecosystems; providing country assessments to inform digital development; and supporting the responsible use of digital technologies. For example, USAID will support the African Union’s drafting of a dedicated protocol on digital trade and e-Commerce within the African Continental Free Trade Area. However, global demand significantly exceeds available funding for the digital programs. In FY2019, for example, USAID provided only \$1.7 million in foreign assistance under the Digital Ecosystem Fund, one component of the strategy.</p>
	<p>DFC’s Roadmap for Impact. The Roadmap proposes catalyzing \$75 billion—\$25 billion in DFC funds plus \$50 billion from the private sectors—over five years (2020-2025) to support development projects in low- and lower middle-income countries. DFC’s investment focus on “Technology and Infrastructure” includes support for open, interoperable, reliable and secure digital infrastructure and internet access. The DFC five-year plan aspires to commit \$5 billion in technology and critical infrastructure investments through 10 major infrastructure projects.</p>
	<p>EXIM’s Program on China and Transformational Exports (PCTE). Established by Congress in December 2019, PCTE directs EXIM to reserve no less than 20% of its total financing authority (\$27 billion out of \$135 billion) “to directly neutralize export subsidies for competing goods and services financed by official export credit, tied aid, or blended financing provided by China or by other covered countries” and “to advance the comparative leadership of the U.S. with respect to China, or support United States innovation, employment, and technological standards, through direct exports” in specific industries. AI and associated technologies are an explicit focus.</p>
	<p>USTDA’s Global Infrastructure Resilience Initiative. The Initiative supports work with emerging markets to plan, sustain, and finance infrastructure to combat external threats. Projects include ICT, remote health technology, emergency response technology, and grid infrastructure technology. USTDA also collaborates with USAID and other agencies to build digital infrastructure and provide technical assistance as part of the DCCP.</p>

Advancing Sustainable Infrastructure in the Indo-Pacific Region, USAID (last accessed Feb. 16, 2021), https://www.usaid.gov/sites/default/files/documents/1861/USAID_ITAN_Fact_Sheet_080719.pdf; Fact Sheet, U.S. Department of State, *U.S.-ASEAN Smart Cities Partnership (USASCEP): Sharing Expertise Between Cities to Benefit the People of ASEAN* (Feb. 12, 2021), <https://www.state.gov/u-s-asean-smart-cities-partnership-usascp-sharing-expertise-between-cities-to-benefit-the-people-of-asean/>; *U.S.-ASEAN Smart Cities Partnership*, USASCP (last accessed Feb. 16, 2021), <https://www.usascp.org/home-page>; *Digital Connectivity and Cybersecurity Partnership (DCCP)*, USAID (Oct. 19, 2020), <https://www.usaid.gov/digital-development/digital-connectivity-cybersecurity-partnership>; Fact Sheet, U.S. Embassy and Consulate in the Republic of Korea, *2020 Indo-Pacific Business Forum Promote Free and Open Indo-Pacific* (Oct. 29, 2020), <https://kr.usembassy.gov/102920-2020-indo-pacific-business-forum-promotes-free-and-open-indo-pacific/>; Press Release, U.S. Embassy Chile, *U.S. Support for Digital Transformation in Latin America and the Caribbean* (Nov. 10, 2020), <https://cl.usembassy.gov/u-s-support-for-digital-transformation-in-latin-america-and-the-caribbean/>; Fact Sheet, U.S. Department of State, *U.S. Support for Digital Transformations in Latin America and the Caribbean* (Nov. 10, 2020), <https://2017-2021.state.gov/u-s-support-for-digital-transformation-in-latin-america-and-the-caribbean//index.html>; USAID has published resources, such as *Reflecting the Past, Shaping the Future: Making AI Work for International Development and Managing Machine Learning Projects in International Development: A Practical Guide*, which are reflective of the Digital Strategy and outline democratic principles in the deployment of those technologies. See *Reflecting the Past, Shaping the Future: Making AI Work for International Development*, USAID (May 21, 2019), <https://www.usaid.gov/digital-development/machine-learning/AI-ML-in-development>; *Managing Machine Learning Projects in International Development: A Practical Guide*, USAID (Jan. 12, 2021), <https://www.usaid.gov/digital-development/managing-machine-learning-projects>; *Digital Strategy 2020-2024*, USAID, (Jun. 2020), <https://www.usaid.gov/usaid-digital-strategy>; *Digital Ecosystem Fund: 2020 Activities*, USAID (Dec. 22, 2020), <https://www.usaid.gov/digital-development/DEF2020>; *Roadmap for Impact*, DFC (last accessed Feb. 16, 2021), <https://www.dfc.gov/roadmap-for-impact>; Fact Sheet, Export-Import Bank of the United States, *Overview: Program on China and Transformational Exports* (last accessed Feb. 16, 2021), <https://www.exim.gov/who-we-serve/external-engagement/china-and-transformational-exports-program/fact-sheet#:~:text=EXIM%20is%20actively%20working%20to,or%20by%20other%20covered%20countries>; Fact Sheet, Export-Import Bank of the United States, *Program on China and Transformational Exports: Supporting Artificial Intelligence* (last accessed Feb. 16, 2021), <https://www.exim.gov/who-we-serve/external-engagement/supporting-us-artificial-intel-exports>; Press Release, Export-Import Bank of the United States, *EXIM Board Unanimously Approves Historic Policy to Support U.S. Exporters Competing with the People's Republic of China* (Dec. 18, 2020), <https://www.exim.gov/news/exim-board-unanimously-approves-historic-policy-support-exporters-competing-peoples-republic>; *Global Infrastructure Resilience Initiative*, U.S. Trade and Development Agency (last accessed Feb. 16, 2020), <https://ustda.gov/call-for-initial-proposals/>; *2021 Indo-Pacific Strategy Fact Sheet*, U.S. Trade and Development Agency (last accessed Feb. 16, 2020), <https://ustda.gov/wp-content/uploads/2021-Indo-Pacific-Strategy-Fact-Sheet.pdf>; Press Release, U.S. Trade and Development Agency, *USTDA Announces 2021 Digital Strategy for the Indo-Pacific* (Oct. 27, 2020), <https://ustda.gov/ustda-announces-2021-digital-strategy-for-the-indo-pacific/>.

Blueprint for Action: Chapter 15 - Endnotes

¹ For more information on the National Technology Strategy, see Chapter 9 of this report.

² See Chapter 9 of this report for more details.

³ See *Interim Report and Third Quarter Recommendations*, NSCAI at 214-218 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

⁴ The Commission supports these efforts and further encourages the U.S. government to engage proactively through a “mosaic approach” to ensure the Emerging Technology Coalition is additive and not duplicative. See *Interim Report and Third Quarter Recommendations*, NSCAI at 185 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

⁵ Key organizations for membership include the OECD, the Council of Europe (through the Ad Hoc Committee on AI), the Freedom Online Coalition, GPAI, and the North Atlantic Treaty Organization. The United States and other core partner states should consider including international organizations as observers (e.g., the World Bank, the International Monetary Fund, the World Intellectual Property Organization, the World Health Organization, the World Trade Organization, and United Nations Educational, Scientific, and Cultural Organization).

⁶ *Recommendation of the Council on Artificial Intelligence*, OECD (May 21, 2019), <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

⁷ The IDDI should mobilize financial resources and technical expertise as the DFI Alliance, a partnership between the U.S. International Development Finance Corporation and development finance institutions (DFIs) of 15 other OECD countries, ushered in response to COVID-19. See *Development Finance Institutions Join Forces to Respond to COVID-19 in Developing Countries*, U.S. International Development Finance Corporation (April 6, 2020), <https://www.dfc.gov/media/press-releases/development-finance-institutions-join-forces-respond-covid-19-developing>.

⁸ The Global Infrastructure Hub has forecasted global telecommunications infrastructure investment need at \$8.9 trillion over the next approximately 20 years, with current trends falling short of the need by \$1 trillion. *Forecasting Infrastructure Investment Needs and Gaps*, Global Infrastructure Hub (2020), <https://outlook.gihub.org/>. The Alliance for Affordable Internet estimates it will cost \$428 billion and up to 10 years to achieve universal connectivity to quality broadband internet. See Maiko Nakagaki, *\$428 Billion Investment Needed to Connect All of Humanity to the Internet by 2030*, Alliance for Affordable Internet (Sept. 17, 2020), <https://a4ai.org/428-billion-investment-needed-to-connect-all-of-humanity-to-the-internet-by-2030/>.

⁹ This could include bilateral development finance institutions (DFIs) in OECD member countries and multilateral DFIs, which are the private-sector arms of multi-state IFIs. See *Development Finance Institutions and Private Sector Development*, OECD (last accessed Jan. 27, 2021), <https://www.oecd.org/development/development-finance-institutions-private-sector-development.htm>.

¹⁰ See the Chapter 7 Blueprint for Action and the Annex of this report containing the abridged version of NSCAI’s Key Considerations for Responsible Development & Fielding of AI. For additional details on the Commission’s recommendation to employ technologies and operational policies that align with privacy preservation, see the section on “Aligning Systems and Uses with American Values and the Rule of Law” in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹¹ *G20 Principles for Quality Infrastructure Investment* (last accessed Jan. 4, 2021), https://www.mof.go.jp/english/international_policy/convention/g20/annex6_1.pdf.

¹² “The Digital Principles were first created in consultation with organizations such as The Bill and Melinda Gates Foundation, the Swedish International Development Agency (SIDA), the UN’s Children’s Fund (UNICEF), UN Development Program (UNDP), the World Bank, and the U.S. Agency for International Development (USAID), and the World Health Organization (WHO).” See *Frequently Asked Questions, Principles for Digital Development* (last accessed Jan. 3, 2021), <https://digitalprinciples.org/about/>.

¹³ *Criteria for Security and Trust in Telecommunications Networks and Services*, CSIS Working Group on Trust and Security in 5G Networks (May 2020), https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/200511_Lewis_5G_v3.pdf.

¹⁴ *Recommendation of the Council on Digital Security of Critical Activities*, OECD (Oct. 12, 2019), <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0456>.

¹⁵ *OECD Guidelines for Multinational Enterprises*, OECD (2011), <http://www.oecd.org/daf/inv/mne/48004323.pdf>.

¹⁶ *Guiding Principles on Business and Human Rights: Implementing the United Nation “Protect, Respect and Remedy” Framework*, United Nations Human Rights Office of the High Commissioner (2011), https://www.ohchr.org/Documents/Publications/GuidingPrinciplesBusinessHR_EN.pdf. To mark the 10th anniversary of the adoption of the UN Guiding Principles on Business & Human Rights (UNGPs), the UN will review existing gaps and develop a roadmap for the next decade. See *UN Guiding Principles: The Next Decade*, Business & Human Rights Resource Centre (last accessed Feb. 8, 2021), <https://www.business-humanrights.org/en/big-issues/un-guiding-principles-on-business-human-rights/un-guiding-principles-the-next-decade/>.

¹⁷ *Development Finance Standards*, OECD (last accessed Jan. 27, 2020), <http://www.oecd.org/dac/financing-sustainable-development/development-finance-standards/>; *DAC and CRS Code Lists*, OECD (last accessed Dec. 28, 2020), <http://www.oecd.org/dac/financing-sustainable-development/development-finance-standards/dacandcrscodelists.htm>.

¹⁸ *Recommendation of the Council on Artificial Intelligence*, OECD (May 21, 2019), <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449>.

¹⁹ See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details, see the *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

²⁰ See Chapter 14 of this report for additional details and recommendations pertaining to end-use and end-user export controls.

²¹ For example, Hillhouse Capital, an Asia-focused private equity firm known for its early investments in Tencent and Baidu, grew from a “boutique hedge fund into a \$60 billion behemoth that’s made prescient bets on stocks, private equity and venture capital.” Hillhouse currently seeks to raise “what would be Asia’s largest U.S. dollar-denominated fund targeting \$13 billion.” See Michael McDonald & Lulu Yilun Chen, *Hillhouse Reloads After Building \$60 Billion Asia Juggernaut*, Bloomberg (April 28, 2020), <https://www.bloomberg.com/news/articles/2020-04-27/yale-s-2-4-billion-profit-machine-hillhouse-ready-to-reload>; Kane Wu & Julie Zhu, *Exclusive: Hillhouse Targets Over \$3 Billion for New Yuan-Denominated Fund: Sources*, Reuters (Sept. 18, 2020), <https://www.reuters.com/article/us-hillhouse-fundraising-exclusive/exclusive-hillhouse-targets-over-3-billion-for-new-yuan-denominated-fund-sources-idUSKBN2690LK>.

²² See Chapters 11, 13, 14, and 16 of this report, along with their associated Blueprints for Action, for recommendations to strengthen public-private partnerships and private-sector investments in the United States.

²³ *Interim Report and Third Quarter Recommendations*, NSCAI at 205 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²⁴ This executive order would build upon Executive Order 13859. Donald J. Trump, *Executive Order on Maintaining American Leadership in Artificial Intelligence*, The White House (Feb. 11, 2019), <https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>; see also *Interim Report and Third Quarter Recommendations*, NSCAI at 207-12 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²⁵ *Interim Report and Third Quarter Recommendations*, NSCAI at 206 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²⁶ See the Chapter 15 Blueprint for Action Annex. For additional recommendations on how NIST can support qualified confidence in AI models and predicted outcomes, see Chapter 7 of this report and its associated Blueprint for Action.

Blueprint for Action: Chapter 15 - Endnotes

²⁷ This includes the American National Standards Institute, the primary industry organization advocating for U.S. companies before international standards bodies, and the International Digital Economy and Telecommunication Advisory Committee (IDET). See *Interim Report and Third Quarter Recommendations*, NSCAI at 208-209 (Oct. 2020); About IDET, U.S. Department of State (last accessed Feb. 11, 2021), <https://www.state.gov/international-digital-economy-and-telecommunication-advisory-committee/about-idet/>.

²⁸ *Interim Report and Third Quarter Recommendations*, NSCAI at 206 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²⁹ *Id.*

³⁰ *Id.*

³¹ The Task Force should also evaluate efforts to promote targeted development priorities. For example, Power Africa is a “U.S. Government-led partnership, coordinated by USAID, that brings together the collective resources of over 150 public- and private-sector partners to double access to electricity in sub-Saharan Africa.” See *Power Africa*, USAID (Nov. 30, 2020), <https://www.usaid.gov/sites/default/files/documents/power-africa-fact-sheet-11-2020.pdf>.

³² This evaluation should include a review of the Defense Production Act (DPA) as a tool for DFC and potentially other agencies to promote the U.S. industrial base, as was done as part of the response to COVID-19. See *Defense Production Act (DPA)*, U.S. International Development Finance Corporation (last accessed Jan. 4, 2021), <https://www.dfc.gov/dpa>.

³³ The Digital Strategy includes several complementary efforts relating to connectivity, cybersecurity, digital finance, inclusion, and other areas, such as the Digital Connectivity and Cybersecurity Partnership, Digital Finance, Digital Inclusion, Geospatial Technology and Analytics, Development Informatics, and Digital Agriculture, among others. Critical components of the Digital Strategy also include catalytic funding provided to missions (Digital Ecosystem Fund) and conducting Digital Ecosystem Country Assessments (DECAs). See *USAID Digital Strategy 2020-2024*, USAID (June 2020), <https://www.usaid.gov/usaid-digital-strategy>. USAID has also published resources which are reflective of the Digital Strategy and outline democratic principles in the deployment of those technologies. See, e.g., Amy Paul, et al., *Reflecting the Past, Shaping the Future: Making AI Work for International Development*, USAID (May 2, 2019), <https://www.usaid.gov/sites/default/files/documents/15396/AI-ML-in-Development.pdf>; *Artificial Intelligence in Global Health: Defining a Collective Path Forward*, USAID (2020), https://www.usaid.gov/sites/default/files/documents/1864/AI-in-Global-Health_webFinal_508.pdf.

³⁴ For example, USAID may want to explore the potential for all its programs to include a minimum threshold of digital programming (e.g., 10% of programmatic efforts include an element of digital development) as a key element of supporting nations on their journey to self-reliance.

³⁵ *Roadmap for Impact*, DFC (last accessed Nov. 25, 2020), <https://www.dfc.gov/roadmap-for-impact> [hereinafter DFC Roadmap for Impact].

³⁶ DFC *Roadmap for Impact* at 6.

³⁷ DFC *Roadmap for Impact* at 57; *Congressional Budget Justification: Fiscal Year 2021*, DFC (last accessed Nov. 25, 2020), https://www.dfc.gov/sites/default/files/media/documents/FY2021_DFC_CBJ-Final-04222020.pdf [hereinafter DFC FY2021 Budget].

³⁸ The China Development Bank and China Exim Bank provide concessional loans, including, for example, a 40-year concessionary loan to Indonesia to fund its U.S. \$5.29 B high-speed railway. The loan provided a 10-year grace period, no guarantees by Indonesia, and local content guarantees. See *China's Belt and Road Initiative in the Global Trade, Investment, and Finance Landscape*, OECD at 18 (2018), <https://www.oecd.org/finance/Chinas-Belt-and-Road-Initiative-in-the-global-trade-investment-and-finance-landscape.pdf>.

³⁹ *Policy Update*, Donor Tracker (July 23, 2020), <https://donortracker.org/policy-updates/european-investment-bank-provide-us84-million-concessional-loan-senegal-support>; *Japanese Concessional ODA Loans*, United Nations (last accessed Feb. 2, 2021), <https://www.un.org/ldcportal/japanese-concessional-oda-loans/>; *Understanding China's Belt and Road Infrastructure Projects in Africa*, Brookings Institution (Sept. 2019), https://www.brookings.edu/wp-content/uploads/2019/09/FP_20190930_china_bri_dollar.pdf.

⁴⁰ DFC has formal relationships with international partners such as Japan and Australia (Japan Bank for International Cooperation, Nippon Export Investment Insurance, Australia Infrastructure Financing Facility), the African Development Bank, and the Inter-American Development Bank. See Testimony by Adam S. Boehler, CEO, U.S. International Development Finance Corporation before the House Appropriations Subcommittee on State and Foreign Operations, and Related Programs (March 4, 2020), <https://www.dfc.gov/testimony-DFC-HAP-03042020>.

⁴¹ Blended finance, according to the OECD, “is the strategic use of development finance for the mobilisation of additional finance towards sustainable development in developing countries.” Blended concessional finance includes the “use of relatively small amounts of concessional donor funds to mitigate specific investment risks and help rebalance risk-reward profiles of pioneering investments that are unable to proceed on strictly commercial terms.” See *Blended Finance*, OECD (last accessed Dec. 28, 2020), <https://www.oecd.org/dac/financing-sustainable-development/blended-finance-principles/>; *Blended Concessional Finance*, International Finance Corporation of the World Bank Group (last accessed Dec. 28, 2020), https://www.ifc.org/wps/wcm/connect/topics_ext_content/ifc_external_corporate_site/bf; see also *Blended Finance*, Convergence (last accessed Dec. 28, 2020), <https://www.convergence.finance/blended-finance>.

⁴² This investment fund could be modelled after the \$100 million Women’s World Banking Asset Management Fund, which received \$25 million from DFC and \$600,000 from USAID. Additionally, DFC put out a call for proposals for fund managers investing in 5G-related companies operated in DFC-eligible emerging market countries in order to invest in open and secure ICT. See *First-of-Its-Kind U.S. Government Blended Finance Fund to Empower Women in Developing Markets*, DFC (June 4, 2019), <https://www.dfc.gov/media/opic-press-releases/first-its-kind-us-government-blended-finance-fund-empower-women>; *Information and Communication Technology Call for Proposals*, DFC (March 2020), https://www.dfc.gov/sites/default/files/media/documents/IFD_5G_CFP_032020.pdf.

⁴³ The \$82 million request was part of the larger \$640.1 million requested to support DDI’s efforts. See *Congressional Budget Justification: Foreign Operations: Appendix 2: FY 2021*, U.S. Department of State at 223 (2020), <https://www.state.gov/wp-content/uploads/2021/01/FY21-CBJ-Appendix-2-FINAL-508-Version.pdf>.

⁴⁴ The Digital Ecosystem Fund “equips the Agency’s Operating Units with catalytic financing to design and implement activities that foster open, inclusive, and secure digital ecosystems. The DEF supports two types of interventions: 1) Emergent opportunities to harness or shape the digital ecosystem in any sector; 2) Strategic initiatives to strengthen or improve the digital ecosystem.” See *Digital Ecosystem Fund: 2020 Activities*, USAID (Dec. 22, 2020), <https://www.usaid.gov/digital-development/DEF2020>.

⁴⁵ DFC’s FY 2021 budget request sought \$700 million in such funds. See DFC FY2021 Budget at 1.

⁴⁶ The U.S. government announced a first-of-its-kind blended finance fund in June 2019. USAID provided \$600,000 in funding and technical assistance and DFC’s predecessor invested \$25 million to support private capital investments in the \$100 million Women’s World Banking Asset Management Fund. See *First-of-Its-Kind U.S. Government Blended Finance Fund to Empower Women in Developing Markets*, DFC (June 4, 2019), <https://www.dfc.gov/media/opic-press-releases/first-its-kind-us-government-blended-finance-fund-empower-women>; Vince Chadwick, *USAID, OPIC Team Up on Women’s Finance in ‘Preview’ of New DFI Era*, Devex (June 5, 2019), <https://www.devex.com/news/usa-id-opic-team-up-on-women-s-finance-in-preview-of-new-dfi-era-95050>.

⁴⁷ See, e.g., DFC Roadmap for Impact at 57.

⁴⁸ DFC FY2021 Budget; DFC Roadmap for Impact.

⁴⁹ DFC Roadmap for Impact at 56.

⁵⁰ See Chapter 14 of this report for recommendations regarding specific end-use controls on high-end AI chips.

⁵¹ 85 Fed. Reg. 43532, *Advanced Surveillance Systems and Other Items of Human Rights Concern*, U.S. Department of Commerce: Bureau of Industry and Security (July 17, 2020), <https://www.federalregister.gov/documents/2020/07/17/2020-15416/advanced-surveillance-systems-and-other-items-of-human-rights-concern>; 85 Fed. Reg. 63007, *Amendment to Licensing Policy for Items Controlled for Crime Control Reasons*, U.S. Department of Commerce: Bureau of Industry and Security (Oct. 6, 2020), <https://www.federalregister.gov/documents/2020/10/06/2020-21815/amendment-to-licensing-policy-for-items-controlled-for-crime-control-reasons>.

Blueprint for Action: Chapter 15 - Endnotes

- ⁵² *Guidance on Implementing the UN Guiding Principles for Transactions Linked to Foreign Government End-Users for Products or Services with Surveillance Capabilities*, U.S. Department of State (last accessed Jan. 4, 2021), <https://www.state.gov/wp-content/uploads/2020/10/DRL-Industry-Guidance-Project-FINAL-1-pager-508-1.pdf>.
- ⁵³ Kara Frederick, *Democracy by Design: An Affirmative Response to the Illiberal Use of Technology for 2021*, Center for a New American Security (Dec. 15, 2020), <https://www.cnas.org/publications/reports/democracy-by-design>; Dahlia Peterson, *Designing Alternatives to China's Repressive Surveillance State*, Center for Security and Emerging Technology (Oct. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-Designing-Alternatives-to-Chinas-Surveillance-State.pdf>.
- ⁵⁴ GPAI was launched in June 2020 to advance “responsible and human-centric” AI consistent with human rights, fundamental freedoms, democratic values, innovation, and economic growth. Current members include Australia, Brazil, Canada, the European Union, France, Germany, India, Italy, Japan, Mexico, the Netherlands, New Zealand, Poland, Singapore, Slovenia, South Korea, Spain, the United Kingdom, and the United States, with the OECD and UNESCO as Permanent Observers. See *UNESCO Joins Global Partnership on Artificial Intelligence as Observer*, UNESCO (Dec. 10, 2020), <https://en.unesco.org/news/unesco-joins-global-partnership-artificial-intelligence-observer>.
- ⁵⁵ GPAI's five working groups (responsible AI, data governance, innovation and commercialization, the future of work, and pandemic response) are supported by research undertaken by two centres of expertise: the Paris-based National Institute for Research in Digital Science and Technology (INRIA) and the Montreal-based International Centre of Expertise in Montreal for the Advancement of Artificial Intelligence (ICEMAI). See *The Global Partnership on Artificial Intelligence Officially Launched*, Montreal International (June 15, 2020), <https://www.montrealinternational.com/en/news/the-global-partnership-on-artificial-intelligence-officially-launched/>; *Launch of the Global Partnership on Artificial Intelligence by 15 Founding Members*, French Ministry for Europe and Foreign Affairs (June 15, 2020), <https://www.diplomatie.gouv.fr/en/french-foreign-policy/digital-diplomacy/news/article/launch-of-the-global-partnership-on-artificial-intelligence-by-15-founding>.
- ⁵⁶ The OECD's work through the Directorate for Science, Technology and Innovation and the AI Policy Observatory is supported by partnerships with governments and research entities, like the German AI Observatory's support of the OECD's effort on AI's impact on the labor market. See *Work, Innovation, Productivity and Skills Programme: Overview*, OECD.AI (last accessed Feb. 1, 2021), <https://oecd.ai/work-innovation-productivity-skills>.
- ⁵⁷ *Going Digital*, OECD (last accessed Feb. 2, 2021), <http://www.oecd.org/going-digital/project/>.
- ⁵⁸ Banff International Research Station (last accessed Jan. 4, 2021), <https://www.birs.ca/>.
- ⁵⁹ AccelNet accelerates network-to-network collaborations by funding the connection (travel, virtual networking, workshops) between international research networks. NSF only funds the U.S. portion and expects international partners to fund their part of the collaboration. In addition to funding connections between existing networks, AccelNet will fund efforts to create and foster nascent networks. See *Accelerating Research Through International Network-to-Network Collaborations (AccelNet)*, NSF (Sept. 21, 2020), https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505584&org=OISE&from=home.
- ⁶⁰ MULTIPLIER sends subject-matter experts to international areas of interest to assess capabilities and gather information for potential NSF joint projects. This new program has already been promising—NSF sent a multidisciplinary team to the Czech Republic and was astounded at their capabilities. NSF is now exploring bilateral collaboration. See *NSF MULTIPLIER: MULTIPLYing Impact Leveraging International Expertise in Research Missions*, NSF (last accessed Dec. 28, 2020), <https://www.nsf.gov/od/oise/multiplier.jsp>; NSCAI staff discussions with NSF staff (Nov. 14, 2020).
- ⁶¹ For example, the Department of Energy's national laboratories may be used to sponsor research with the recommended dedicated funding.
- ⁶² The shared research resource can help prevent bottlenecks due to limited compute resources. This effort may also be, if appropriate, part of an expansion of the National AI Research Resource delineated in Chapter 9 of this report.

⁶³ NSCAI recommends each member dedicate funding to support research efforts. However, MAIRI will also serve as a location for research to gather for dialogues, workshops, and mentorships. Based on similar international research institutes, MAIRI members should consider providing the equivalent of \$100K-\$250K per year to cover the travel, accommodations, and per diem of around 80 researchers to MAIRI to facilitate communications and interactions between researchers.

⁶⁴ In 2017, the government of Alberta (Canada), Canadian Natural Science and Engineering Research Council, U.S. National Science Foundation, and Mexico's Consejo Nacional de Cinco y Tecnología invested \$12.5 million over the next five years. See *Research Station Gets \$12.5M to Bring Scientists and Mathematicians to Banff*, CBC News (Feb. 10, 2017), <https://www.cbc.ca/news/canada/calgary/banff-international-research-station-math-science-funding-1.3977703>. For the 2012-2017 period, BIRS received \$10.3 million, of which \$3.68 million was from NSF. For the 2006-2011 period, BIRS received \$9.3 million, \$3.1 million of which was from NSF. See *Organization: Banff International Research Station*, Research Money (last accessed Jan. 4, 2021), <https://researchmoneyinc.com/organization/banff-international-research-station/>.

⁶⁵ The National AI Research Institutes provide an important example of the power of leveraging R&D cooperation across the U.S. interagency. The new NSF and DOE centers for quantum information science are also a powerful example of the benefits of dedicated resourcing and prioritization across two U.S. government research entities. See Andrea Peterson, *NSF and DOE Support Research Priorities with Spate of New Center Awards*, American Institute of Physics (Sept. 16, 2020), <https://www.aip.org/fyi/2020/nsf-and-doe-support-research-priorities-spate-new-center-awards>.

⁶⁶ For more information, see the Chapter 10 Blueprint for Action.

⁶⁷ J-1 visas are used by academic employers like universities and research institutions to sponsor foreign-born academics, interns, trainees, and researchers to work for several months to five years in the United States. J-1 visa holders are not allowed to renew their visas and must wait one to two years before they can apply for a different immigration status. O-1 visas are for individuals who can provide extensive evidence that they have “extraordinary ability in the sciences, arts, education, business, or athletics.” These visas last for up to three years with indefinite renewals but have been used minimally to attract experts in S&T due to restrictive policy guidance. For a discussion of immigration and visa programs to attract scientists and researchers to the United States, see generally Zachary Arnold, et al., *Immigration Policy and the U.S. AI Sector*, CSET (Sept. 2019), <https://cset.georgetown.edu/research/immigration-policy-and-the-u-s-ai-sector/>.

⁶⁸ See also Chapter 1 of this report on Malign Information Operations.

⁶⁹ *Secretary Pompeo Approves New Cyberspace Security and Emerging Technologies Bureau*, U.S. Department of State (Jan. 7, 2021), <https://2017-2021.state.gov/secretary-pompeo-approves-new-cyberspace-security-and-emerging-technologies-bureau/index.html>.

⁷⁰ For additional details regarding the function and need for the CSET bureau, see *Second Quarter Recommendations*, NSCAI at 88-89 (July 2020), <https://www.nscai.gov/previous-reports/>.

⁷¹ These elements should include key, technology-related functions of the proposed CSET Bureau; the Bureau of Oceans, Environment and Science (OES); the Office of the Science and Technology Adviser to the Secretary (STAS); the Coordinator for Cyber Issues (S/CCI); the Bureau for Economic and Business Affairs (EB); the Bureau for Democracy, Human Rights, and Labor (DRL); and the Center for Analytics.

⁷² *Key Topics*, Office of the Science and Technology Adviser at the U.S. Department of State (last accessed Dec. 15, 2020), <https://www.state.gov/key-topics-office-of-the-science-and-technology-advisor/>.

⁷³ *Diplomats in Residence*, U.S. Department of State (last accessed Feb. 2, 2021), <https://careers.state.gov/connect/dir/>.

⁷⁴ See, e.g., *Managing American Spaces*, U.S. Department of State (last accessed Feb. 1, 2021), <https://americanspaces.state.gov/>; *TechCamp*, U.S. Department of State (last accessed Feb. 1, 2021), <https://techcamp.america.gov/>; *Program Description*, World Learning (last accessed Feb. 1, 2021), <https://www.worldlearning.org/program/u-s-speaker-program/>.


⁷⁵ See *About GIST*, GIST (last accessed Feb. 1, 2021), <https://www.gistnetwork.org/about>.

Chapter 15: A Favorable International Technology Order

Annex: Emerging Technology Coalition

This Annex provides a framework and overarching agenda for global cooperation on artificial intelligence (AI) and emerging technologies. It includes guidance on concrete, operational projects; applications; and implementation mechanisms for collaborative AI work across seven critical areas. Collaborative work in these areas will serve to further AI consistent with democratic values and strengthen the ties that connect the United States with its allies and partners. This Annex is intended to provide guidance to the Emerging Technology Coalition (ETC) and may assist officials in prioritizing bilateral and multilateral collaborative efforts outside the context, to include engagement with multilateral initiatives across the AI landscape, as reflected in the Key Multilateral Technology Initiatives Table of the Chapter 15 Blueprint for Action.


Seven Critical Areas for International Cooperation



STANDARDS & NORMS

Advance norms that uphold democratic values


Support secure, reliable, trusted technologies at international technical standards bodies



FACILITATE DATA SHARING


Address legal and regulatory barriers

Explore agreements to enable data sharing and free flow of data



DEVELOP AI-RELATED TALENT

Analyze labor market challenges; harmonize skills and certification requirements; and increase talent exchanges, joint training, and workforce development initiatives



PROMOTE & PROTECT INNOVATION

Export controls
Investment screening
Supply chain assurance
Emerging tech investment
Trade policy
Intellectual property
Research and cyber protections



JOINT R&D ON AI AND DIGITAL INFRASTRUCTURE

Promote collaborative R&D that advances shared interests and benefits humanity



PROMOTE DEMOCRACY, HUMAN RIGHTS, AND THE RULE OF LAW

Pursue joint efforts to counter censorship, malign information operations, human trafficking, and illiberal uses of surveillance technologies



LAUNCH INTERNATIONAL DIGITAL DEMOCRACY INITIATIVE

Coordinate international foreign assistance, policy and technical guidance, and development aid and financing

Critical Area #1 – Developing and Operationalizing Standards and Norms

- **Objectives:**

- o Advance common, democratic norms and values to govern and guide responsible artificial intelligence (AI) and the research, development, and application of emerging technologies globally.¹
- o Promote international AI norms and standards that uphold democratic values, building on guiding documents such as the Organization for Economic Co-Operation and Development (OECD) AI Principles and efforts to operationalize principles, as reflected in the Commission's *Key Considerations for Responsible Development & Fielding of Artificial Intelligence*.²
- o Coordinate positions taken by partner states' governmental delegations and, where appropriate, national (non-government) standards institutes accredited to international technical standards organizations to support development of secure, reliable, and trusted technologies and to ensure ethical and technical integrity, endorse standards that comport with democratic values, and maintain the neutrality of these organizations.³

- **Priority areas for coordination.** The ETC should seek to align with allies and partners and prioritize efforts in development of international technical standards in the following priority areas:

- o Safety and reliability;
- o Privacy-enhancing technologies, including privacy-preserving machine learning (PPML), allied cryptographic code, and other privacy-enhancing technologies;
- o Data sharing, labelling, and related documentation for data, models, and systems;
- o Assessing system performance and characterizing blind spots per shared values (including fairness, interpretability, reliability, and secure use of AI technologies as part of integrated systems);
- o Robustness to ensure models are resilient to adversarial examples and model inversion, while red-teaming with allies on competitors' attempts to undermine AI-enabled systems;
- o Trust in human-machine teaming and development of common standards and benchmarks to assess risks in settings of varying complexity and uncertainty;
- o Traceability, focused on audit trail requirements per mission needs for high-stakes AI systems including safety-critical applications; and
- o Interoperability, including benchmarks that assess reliability of produced model explanations.

- **Mechanisms for Coordination.**

- o In addition to coordinating efforts through the ETC, the U.S. government, led by the Department of State and in coordination with the National Institute of Standards and Technology (NIST) and its AI Standards Coordinator, should engage with democratic nations to align positions on standards that are critical to mutual security and defense.

- The Department of State, as the Commission has recommended, is in the process of placing regional technology officers in major foreign technology hubs.⁴ This development will facilitate diplomatic efforts toward coordinating positions with allies and partners.⁵
- As the Commission recommended in its *Third Quarter Recommendations*⁵ and elsewhere in this Report, NIST and other agencies should consider the Commission's *Key Considerations for Responsible Development and Fielding of AI*. The *Key Considerations* include operational guidance on standards critical to responsible AI and national security, including for technical standards on testing and evaluation, verification and validation (TEVV).⁷
- Coordination on technical standards should include the work of international standards organizations as well as coordinated work on operationalizing AI norms and principles. The Global Partnership on AI's (GPAI) Data Governance Working Group can provide particularly salient best practices for engineering and implementing data-sharing, pooling, and collecting initiatives.⁸ The OECD is also a critical forum for technical standards and guidelines, particularly in data sharing and responsible AI.⁹

Critical Area #2 – Joint Research and Development on AI and Digital Infrastructure

• Objectives:

- Identify areas of shared interest conducive to collaborative R&D—such as privacy-enhancing technologies, small data approaches to AI, next-generation materials, prototyping, and high-performance computing (HPC)—for which there are existing gaps and identify ways to share resources to pursue R&D in those areas.
- Develop mechanisms to facilitate fundamental and applied R&D projects that involve collaboration among nations, industry partners, and researchers.
 - Projects may also include secure cloud frameworks, sharing best practices on TEVV, innovative funding models, international test beds to develop pre-commercial technologies, and leveraging the Commission's proposed Multilateral AI Research Institute for coordination (see the Chapter 15 Blueprint for Action).
- Pursue collaborative, coordinated efforts to develop and deploy AI applications to benefit humanity at large in areas of global concern such as those embodied in UN Sustainable Development Goals.

• Priority areas for collaborative R&D—advancing AI technology.

- *Development of privacy-preserving technology*, such as homomorphic encryption and differential privacy techniques,¹⁰ to facilitate cross-border AI applications, data sharing, and cooperative efforts.¹¹
- *Continuous development and adaptation of TEVV systems* to strengthen the development of trustworthy, robust AI is critical to advancing the interests of democratic nations and to understanding how AI systems perform in multi-agent/adversarial contexts.¹² Collaboration in this area will contribute to understanding differences among allies on policy, metrics, standards, and requirements while creating stronger connections for all users in a full-cycle approach.

- o *Development of AI for modeling, simulation, and design* to provide researchers with a larger scope of AI-ready data sets.¹³
- o *Development of one- and few-shot learning algorithms*¹⁴—algorithms that rely on less data—to facilitate future joint R&D and data sharing and improve context-specific interoperability.¹⁵
- o *Development of robust allied AI* to reduce vulnerabilities of allied AI systems and training data to adversarial attacks.¹⁶
- o *Achieving context-specific interoperability of AI systems* necessary for cross-border AI applications, with a focus on how systems integrate particular AI/ML components.¹⁷ The potential for AI to increase speed of operations will require allies and partners to stress-test decision-making procedures and communications protocols to ensure interoperability. Interoperability of AI systems is already an issue at the forefront of defense cooperation and will only grow in importance as technology matures.
- o *Development of AI to secure and improve resiliency of supply chains* to protect AI-component supply chains while promoting domestic and allied innovation and to apply AI to improve auditing, mapping, and securing supply chains while ensuring resilience to shocks. Given the inherently cross-border nature of supply chains and their critical role in the international economy around AI and advanced technology, this is a natural area for the United States to work collaboratively with like-minded nations.
- o *Additional critical AI research areas* including novel machine learning (ML) directions, complex multi-agent scenarios, advanced scene understanding, AI system risk assessment, enhanced human-AI interaction and teaming, and autonomous AI systems.¹⁸
- **Priority areas for collaborative R&D—AI to benefit humanity.** The potential for AI to assist the global community in improving the human condition is immense. Priority areas for international collaboration should include the following:
 - o *Environment and climate.* Recognizing the growing view that environmental degradation and climate change represent imperatives for national and international security, the international community must work collaboratively to develop AI-based solutions to address common climate, environmental, and energy challenges.
 - Collaborative initiatives such as the following serve as models for future international efforts:
 - The Partnership between Cross Section Evaluation Working Group and OECD's Nuclear Energy Agency's (NEA) Working Party on International Nuclear Data Evaluation Co-operation on International Criticality Safety Benchmark Evaluation Project¹⁹;
 - GEOTHERMICA, a collaboration among 12 European countries and the United States to fund AI-specific research on geothermal R&D²⁰; and
 - The International Partnership for Hydrogen & Fuel Cells in the Economy (IPHE),²¹ an intergovernmental partnership to facilitate and accelerate transition to clean and efficient energy with the support of AI and ML research.

- o *Health, including pandemic detection and response.* The COVID-19 crisis has made clear the need for global collaboration and the potential for AI-enabled solutions.
 - *Smart disease monitoring.* The Commission has recommended global cooperation on smart disease monitoring.²² Such a global initiative, for example, could seek to combine existing data on zoological spills with open source health-related data to create shared, predictive, global disease monitoring models (see Chapter 16 of this report and its associated Blueprint for Action).²³
 - *Pandemic preparedness, vaccine development, and syndromic surveillance.* Efforts in this space²⁴ include:
 - Development and coordination on international norms and standards to govern use and sharing of international health data, protecting privacy while ensuring timely accessibility of data;
 - Development of privacy standards for genomic data sets;
 - Increased international cooperation in the COVID-19 High Performance Computing Consortium (potentially through GPAI); and
 - Facilitation of international cooperation with DARPA's work on creating the infrastructure and protocols for data sharing and collaboration at the point of experimentation for drug discovery.
 - *Initiatives to enable long-term quality of life.* Collaboration with allies and partners can facilitate the Commission-recommended focus on harnessing AI to help the elderly live independently longer, assist in managing health and daily tasks, and improve the quality of life, particularly through the application of AI to biomedicine.²⁵
 - The National Nanotechnology Initiative's U.S.-EU Communities of Research,²⁶ along with various national-level efforts by partner nations,²⁷ should serve as models for larger-scale international collaboration.
- o *Food security.* The United States emphasizes agriculture-led growth, resilience, nutrition, and water security, sanitation, and hygiene in its foreign assistance programs. Enhancing the security of water and food of partner nations is needed to disrupt the vicious cycle of poverty, hunger, and conflict.²⁸
 - Agricultural sectors are increasing the use of data-driven technologies such as robotics, satellites, GPS, and drones. Significant data sets are being generated about crop growth, soil characteristics, and weather conditions.²⁹
 - AI and ML-based algorithms can amplify the data sets and hardware to improve real-time monitoring and analysis of agricultural and distribution processes. This can improve efforts to assess needs, enhance productivity and security, and build local capacities and productivity while minimizing environmental impact.
- o *Disaster relief.* AI-enabled technologies are being used to address a range of disaster scenarios, and further work on an international basis should be explored. The World Economic Forum is among those groups calling for greater international collaboration in order to realize the benefits of AI to specifically include the area of disaster relief.³⁰

- Climate- and weather-related disasters like hurricanes, wildfires, and flooding are on the rise, and AI is already being applied to mitigate the effects by locating survivors using unmanned aerial vehicles, removing debris after a disaster, deploying robots to communicate with victims, employing edge technology to obtain the most up-to-date data, running predictive simulations, and leveraging social media reporting.
 - The DoD is partnering with DoE and Microsoft to develop “deep-learning artificial intelligence algorithms to provide near-real-time data to improve the decision-making of first responders engaged in natural disasters and humanitarian assistance efforts,”³¹ and countries like Singapore are working with the Joint Artificial Intelligence Center (JAIC) on this particular National Mission Initiative.³²
 - The United States can work with its allies and partners to effectively predict, model, prepare for, and respond to disasters, as the Commission recommends in Chapter 11 of this report.³³
 - *Civilian space cooperation.* The United States and other space agencies employ AI to tackle a range of space missions—including for visualization of space objects and situational awareness, tracking space debris for satellite collision avoidance, roving the lunar surface, deep-space exploration with autonomous systems, and detection of asteroids that could threaten Earth.³⁴
 - For safe satellite navigation around space debris, the European Space Agency (ESA) and the UK Space Agency both have AI initiatives underway, which suggests potential opportunities for closer U.S. collaboration.³⁵ ESA has already established a partnership with Stanford.³⁶
 - India is also building its space program and has deployed an AI-powered Moon rover.³⁷
 - Russia and China appear to be developing technological solutions to the space collision problem, which could present an area for mutually beneficial cooperation. AI-enabled robotic assistants are also being developed for the International Space Station.³⁸
- **Methods to implement collaborative R&D.** The ETC should also explore vehicles to enable R&D collaboration among government partners and non-governmental organizations.
 - Collaboration must include not only government-to-government efforts, but also methods to partner with researchers at academic research centers and in the private sector.
 - Existing science and technology (S&T) agreements between governments may provide the legal foundation for cooperation, but details will depend on the arrangement at issue.
 - The ETC should prioritize approaches that would facilitate the pooling of resources, reduce redundancies, and support development and socialization of best practices.
 - In addition, the ETC should examine challenges to cross-border, collaborative R&D—such as those around data privacy and data sharing between the United States and European Union—and explore solutions to overcome those challenges.

- Potential methods to implement and further collaborative R&D include:
 - Establishment of the *Multilateral AI Research Institute (MAIRI)*. Proposed by the Commission in this report, MAIRI will serve as a center for multilateral research to coordinate joint efforts to develop technologies and align norms that advance responsible, human-centric, and privacy-preserving AI/ML that better societies.
 - Prioritization of R&D work of the *Global Partnership on AI (GPAI)*. The ETC should leverage existing frameworks wherever possible, and GPAI, supported currently by Canada's International Center of Expertise in Montréal for the Advancement of Artificial Intelligence (ICEMAI) and France's National Institute for Research in Digital Science and Technology (INRIA),³⁹ is among the most promising multilateral, multi-stakeholder initiatives to pursue collaborative R&D and advance AI technology for common causes. The Commission has proposed a greater role for U.S. researchers through a U.S.-based Center of Expertise, leveraging the NSF National AI Research Institutes.⁴⁰
 - Creation of a *joint emerging-tech investment consortium*. Modeled on In-Q-Tel, the consortium would spur investment by the United States and foreign partners in early-stage companies to further development of AI technology that advances and/or protects democratic values. The effort would benefit the United States and its allies and partners through a cross-border platform to engage with startups and entrepreneurs in the AI and emerging-tech space.
 - Within the U.S. government, this effort should draw on State Department's Regional Technology Officers, the Foreign Commercial Service, and USAID missions to identify R&D and prototypes to advance U.S. diplomatic, development, and commercial interests.
 - Launching *multilateral innovation prize competitions*. Modeled on Defense Advanced Research Projects Agency (DARPA) Challenges and XPRIZE Foundation competitions, international innovation prize competitions sponsored by two or more governments would incentivize R&D in fundamental AI or around specific applications necessary for national security and help to pool resources and talent with allies and industry.⁴¹
 - Fostering *allied research at U.S. national labs*. The ETC should consider recommendations for increasing research by allies (potentially a subset of the ETC membership) at U.S. national labs on sensitive topics. Although there are limitations on U.S. national labs to allow foreign researchers, domestically housed research efforts would limit concerns around cross-border data-sharing and cybersecurity and could prove fruitful in R&D necessary for defense and security applications.
 - Development of an *R&D matching platform and a global horizon-scanning capability*. The platform would connect researchers and their projects with funders and partners (governments, philanthropists, venture capitalists, companies, research institutions), providing the U.S. Government with increased visibility into research trends. A horizon-scanning capability of global R&D would complement these efforts and draw on open-source data to give policymakers greater understanding of relevant discoveries and key trends in the field.⁴²
 - Development of an *international test bed for TEVV*. An international test bed for TEVV could be modeled on the National AI Research Institutes⁴³ or the Commission-recommended creation of a NIST-sponsored third-party testing center,⁴⁴ but with a cross-border focus, as well as on the AI4EU project.⁴⁵

- o *Improved collaboration* between centers of excellence, research institutes, and industry consortia through additional coordination by partner governments. This concept would leverage existing and soon-to-be-established centers like the MAIRI, European Union Centres of Excellence, European AI-related Digital Innovation Hubs, the U.S. National AI Research Institutes program, the General Services Administration's AI Center of Excellence (in partnership with the Departments of Agriculture and Health and Human Services as well as the JAIC), the Alan Turing Institute in the U.K., and in Canada the Montreal Institute for Learning Algorithms (Mila),⁴⁶ the Alberta Machine Intelligence Institute (Amii),⁴⁷ and the Vector Institute for Artificial Intelligence.⁴⁸ On the U.S. side, this could involve building on industry and academic efforts like the Stanford Institute for Human-Centered Artificial Intelligence.⁴⁹
- o Fostering of *binational R&D foundations*. ETC members may consider developing targeted, binational R&D efforts modeled on the unique binational foundations that facilitate U.S.-Israel and U.S.-India R&D on cutting edge issues.⁵⁰ These can serve as models for other allies and partners to convene international researchers.

Critical Area #3 – Promoting Democracy, Human Rights, and the Rule of Law

- **Objectives:**

- o Collaborative, coordinated efforts to counter anti-democratic uses of AI and emerging technologies through coordinated policy, regulatory alignment (such as end-user export restrictions), and technology deployment.
- o Potential priorities include countering censorship, countering malign information operations, and promoting democratic models of surveillance technology, although the ETC should explore a range of potential applications.
- o Furthering these normative priorities will build on implementation methods addressed in other Critical Areas, especially #2 (joint R&D), #5 (protecting and promoting innovation), and #7 (the International Digital Democracy Initiative).

- **Countering censorship and authoritarian uses of technology**

- o The ETC should explore efforts to use AI and associated technologies to further internet freedom and counter censorship across the world. This work should be designed to implement the principles adopted in November 2020 by the Freedom Online Coalition, a partnership of 32 governments aligned around promoting human rights and individual freedom.⁵¹
- o The United States should leverage the Open Technology Fund, created by the FY 2021 NDAA, to support this effort, as well as related efforts by the Bureau of Democracy, Human Rights, and Labor (DRL) at the Department of State.⁵²
- o The ETC should coordinate efforts in this space with the Council of Europe's Ad Hoc Committee on AI, established in November 2019 to focus on development, design, and application of AI in areas of human rights, democracy, and the rule of law.⁵³
- o To promote private-sector conduct that comports with shared democratic values, the ETC should develop a proposal for end-user controls that would disincentivize private companies from exporting AI and associated technologies that may be used to suppress and violate human and civil rights.⁵⁴

- **Countering Malign Information Operations**

- o Malign information operations present a growing international challenge that is compounded by the use of AI/ML technologies.
- o This ETC should examine coordinated efforts (outside of the intelligence space) to counter disinformation and other information operations. Joint efforts include detecting, moderating, identifying, and classifying malign information, developing standards and best practices, and training experts.
- o The Commission recommends creation of an International Task Force to Counter and Compete Against Disinformation.⁵⁵ An International Task Force to Counter and Compete Against Disinformation (ITF-CCAD) could be established as a joint project between the United States and multiple countries, as well as the EU and NATO, to further joint efforts to enable content moderation and detection of disinformation, develop standards for identifying and classifying misinformation and disinformation (to include deepfake detection), and share best practices and lessons learned with allies. The private sector, academia, and civil society organizations would be important partners in this effort.
- o The ITF-CCAD should draw best practices from, and should work in coordination with, the Global Internet Forum to Combat Terrorism,⁵⁶ along with efforts of the Department of State's Global Engagement Center's (GEC) Technology Engagement Team (TET); the Federal Bureau of Investigation's Foreign Influence Task Force (FITF); the European External Action Services' Strategic Communication Task Force; the EU "Team Europe" initiative; and the NATO/StratCom Center of Excellence. IFT-CCAD should additionally prioritize stress-testing rapid-response mechanisms and look to fund open-source research.
- o It should explore generating best practices for non-tech solutions, such as media literacy, free press,⁵⁷ and civic engagement initiatives, drawing on notable work by the Center for Strategic and International Studies' Defending Democratic Institutions project and the German Marshall Fund's Alliance for Securing Democracy.

- **Surveillance technology that comports with democratic values.**

- o The ETC should dedicate a multifaceted effort to promoting surveillance technology that supports democratic values.⁵⁸ In particular, the effort should focus on (a) promoting technology that delivers a degree of protection for individual privacy and for civil rights and civil liberties and limits the use of data collected or combined in ways that enable re-identification, and (b) countering the global deployment of surveillance technology used to undermine democratic values and individual rights.
- o Doing so will require coordinated R&D, messaging, and development assistance strategies to support democratic alternatives to technology manufactured in China.⁵⁹
- o Fostering the R&D necessary to provide alternatives will require public-private coordination or partnerships at an international level (see Critical Area #2 for potential mechanisms).
 - Potential stakeholders for such a project include NSF, the National AI Research Institutes, DARPA, NIST, various EU Centres of Excellence, research institutions (such as the Johns Hopkins University Applied Physics Laboratory, the

Massachusetts Institute of Technology Computer Science & AI Laboratory, and the Stanford Institute for Human-Centered AI), GPAI, and non-governmental organizations such as OpenMined.

Critical Area #4 – Exploring Ways to Facilitate Data Sharing

- **Objectives:**

- o Address legal and regulatory barriers to international collaborative work; explore bilateral and multilateral, general and specific approaches to enable data sharing, pooling, and storing consistent with privacy, security, and other fundamental values, including the viability of a Data Free Flow with Trust Agreement.

- **Methods to implement coordinated approaches to data sharing.**

- o *Development of shared data environments.* Development of pooled data storage centers, computational environments, and cloud and edge computing facilities to pool data from different sources for free use by credentialed researchers. An approach like this would prove particularly beneficial to improve data sharing among members to the Five Eyes alliance.
- o *Agreement on foundational data documentation, labelling, archiving, and data organization frameworks at international organizations.* Data agreements among members of alliances (such as NATO) or other international organizations would facilitate support to collaborative R&D endeavors; for example, ongoing efforts at the Organisation of Economic Co-operation and Development (OECD) AI Policy Observatory and Global Partnership on Artificial Intelligence (GPAI).
- o *Agreements to share specific data sets with specified foreign partners.* Narrower in scope than the above two approaches, an agreement of this kind would allow researchers from different countries to access the same data sets for their respective projects. For example, in the context of COVID-19 and health care,⁶⁰ countries would need to address data labelling, data storage, data anonymization, data security, and other issues on a joint basis or through a pilot project.
 - This type of effort could also include joint projects with allies to anonymize⁶¹ high-impact data sets for specific research or initiatives, such as National Institutes of Health data sets and data sets maintained, for various purposes, by DOE, the U.S. Agency for International Development (USAID), the U.S. Food and Drug Administration, DARPA, IARPA, and the Department of State's Center for Analytics.
 - Diplomatic effort is needed to resolve divergent views over what constitutes anonymized data, consent, and matters of public interest.
- o *Ad hoc data sharing arrangements on bilateral or multilateral bases.* The ETC should explore the willingness of strategic allies and partners to engage in targeted, non-treaty data-sharing arrangements. "Innovation sandbox" arrangements may be designed to facilitate specific challenges across all domains—security, health, disinformation, environmental resilience, and so on.
- o *A multilateral data-sharing agreement founded on trust.* The ETC should lead an effort to create a formal, potentially treaty-based approach to data sharing, pooling, and storing with like-minded governments modeled on the data free flow with trust (DFFT) concept introduced by Japan at the June 2019 G20 Summit.

DFFT would permit the free flow of data between authorized parties upon meeting specific standards, including intellectual property (IP), privacy, and cybersecurity protections.⁶² The European Commission endorsed the DFFT concept in December 2020.⁶³

- A *general* DFFT would require significant consideration of data protection, IP protection, privacy shield, and trade issues, both for the United States domestically and for foreign partners.
 - A *specific* DFFT, on the other hand, focused on the free flow of data for particular purposes—such as facilitating pandemic response efforts—would have a greater chance of success and could be a model for targeted data-sharing arrangements in other areas of shared interest.
- o *Development of a secure AI research resource infrastructure.* A secure, cloud-based infrastructure would provide researchers from partnered and allied countries access to compute resources, diverse data sets, and controlled environments to enable testing, for example of privacy-preserving ML techniques. Participating like-minded governments would agree to and comply with common technical standards and norms⁶⁴ and risk-based frameworks that ensure privacy, security, reliability, respect for the rule of law, and other appropriate parameters.
- Such an infrastructure could be developed bilaterally or multilaterally and could be a priority effort of the Commission's proposed MAIRI. Research and academic institutions could support the MAIRI effort with appropriate technical and implementation assistance, while GPAI's Data Governance Working Group could support the development and utilization of engineering best practices.

Critical Area #5 – Promoting and Protecting Innovation

• Objectives:

- o Develop an allied strategy to align and develop regulatory and legal regimes in areas critical to fostering domestic and international innovation. These areas include export controls, investment screening, supply chain assurance, emerging technology investment, trade policy, IP, technology transfer, and research protection.
- o Achieving such a strategy will require an integrated approach among allies and partners, leveraging our full technology toolkit, upgrading capabilities and, where necessary, developing new ones to counter threats. These efforts will require a coordinated strategic coordination plan to raise allied public awareness on issues such as technology-transfer risks.

• Export Controls and Investment Screening

- o The ETC should explore coordinated approaches to export controls and investment screening. Cooperation in these areas is critical to ensure that like-minded nations have the authority to unilaterally institute export controls and block predatory investments that present risks to national and international security.
- o The Commission recommends in Chapter 14 of this report that the United States engage with allies and partners on legal reforms to (a) implement a coordinated approach to AI-related export controls and (b) enhance investment screening procedures and enforcement.⁶⁵

- o The Commission has also recommended as part of Chapter 15's Blueprint for Action that the United States should engage with allies and partners to align policy guidance on exports as part of the International Digital Democracy Initiative (IDDI) to promote technologies that comport with shared values and support free and open societies.⁶⁶
- o As detailed in Chapter 14 of this report, export control priorities should include targeted, high-end semiconductor manufacturing equipment (SME) components needed to produce chips at the 16nm node and below. Additionally, states should explore implementing targeted end-use and end-user controls on specific high-end, AI-specialized chips to prevent their use in human rights violations.⁶⁷
- o Consideration should be given to appropriate economic incentives to support alignment on export control and investment screening.
- o The ETC should also pursue robust collaboration on foreign S&T and investment flow monitoring—to include open-source intelligence—to utilize early warning indicators related to strategic acquisition risks. Further, ETC partners must share best practices to monitor smaller transactions that attempt to skirt existing controls.

• Supply Chain Assurance

- o Leadership of the United States and its allies and partners in emerging technologies is dependent on components sourced from strategic competitors or regions with significant geopolitical risk.
- o The semiconductor manufacturing industry is a prime example of an industry that is critical to U.S. and allied security, but which is heavily concentrated in specific geographic regions and is therefore susceptible to supply chain shocks, particularly in the event of a crisis.⁶⁸
- o The ETC should conduct a supply chain assessment and make recommendations on integrated, multilateral approaches to coordinating critical technology components to enhance international security while reducing collective dependence on strategic competitors.⁶⁹
- o The ETC should also develop a strategic plan to fund key choke point technologies and next-generation materials, approaches, and prototyping capabilities at discovery, manufacturing, and applied scales.⁷⁰

• Emerging Technology Investments

- o Likewise, investments in emerging technologies require coordinated action. 5G presents a test case for the challenges of international and multilateral coordination. The United States and partners have cooperated on developing alternatives to Chinese 5G infrastructure multilaterally and bilaterally.
- o The Commission offered recommendations regarding steps to promote domestic development of 5G technology in its *First Quarter Recommendations* and urged the United States to continue to work closely with key allies and partners on both constructive 5G technical solutions, and to ensure that global 5G networks are safe and secure.⁷¹ Chapter 16 of this report details steps to promote domestic development of biotechnology, 5G, quantum computing, autonomy and robotics, advanced manufacturing, and energy systems, while Chapter 13 of this report details steps to cultivate domestic innovation in microelectronics research and manufacturing.⁷² The United States may engage key allies and partners on these technologies.

- o The ETC can serve as a forum to explore these issues in a coordinated manner.

- **Trade Policy**

- o Trade policy is a key lever for the United States and foreign governments to promote an innovation environment. The ETC should consider coordinated approaches to trade policy to further innovation and strengthen national and international security.

- **Intellectual Property**

- o IP rights and regimes are critical to innovation in AI and emerging technologies. The ETC should explore coordinated approaches to IP that could inform a mutual agenda with the World Intellectual Property Organization's (WIPO) Conversation on AI and Intellectual Property, IP5,⁷³ and forums with broader mandates.
- o Coordination on assistance to nations in developing strong and aligned IP regimes. The ETC can assist the United States and partners in prioritizing assistance to nations in improving their IP regimes to help facilitate innovation while deterring IP theft. A more focused approach, through IP5 and WIPO, may prove more impactful in scope and could help to harmonize efforts to shore up IP with respect to identifiable international challenges.
 - The United States should engage with key allies and partners to align on critical aspects of IP, including patent eligibility for AI and associated technologies, countering China's narrative on winning the innovation competition, IP contractual ecosystem impediments to international collaboration, IP protections for data, and the over-declaration of "standard essential" patents and other efforts to efforts by countries to exploit standards-setting and licensing processes.⁷⁴
 - These are among a set of 10 critical IP considerations that the Commission proposes to guide U.S. efforts to reform IP policies and establish new IP regimes for AI and critical emerging technologies in order to protect and promote national security, innovation, and technology competitiveness.⁷⁵
- o Coordinated efforts to stop IP theft and counter cyber espionage. IP theft remains a global concern. With a goal of protecting the economic viability of AI innovation and emerging technologies, the ETC should identify methods to strengthen the international framework for addressing the export of counterfeit goods, theft of IP technology, forced technology transfers of foreign innovation, and cyber espionage.

- **Research and Cyber Protections**

- o Promoting multilateral responses to research integrity and security. As the Commission has proposed, the United States should coordinate action with allies and partners in developing multilateral responses to challenges to research integrity and security posed by PLA-affiliated individuals and entities and to promote a commitment to open fundamental research.⁷⁶
 - A public-private research security clearinghouse that enables sharing of open-source information, data-driven assessments, decision-support resources, and education and training resources could strengthen this effort.⁷⁷
- o Promoting multilateral efforts to mitigate proliferating cyber vulnerabilities and develop AI-enabled defenses against cyber attacks. As the Commission has proposed, the United States must prepare for AI-enabled cyber conflict. The United States should explore coordinating and joint efforts with key allies and partners.⁷⁸

Critical Area #6 – Developing AI-Related Talent

- **Objectives:**

- o Cooperative efforts to enable government, military, academic, and private-sector talent exchanges and address challenges posed by immigration and visa restrictions; development of joint AI and digital training and workforce-development programs.

- **Methods for furthering talent development globally. The ETC should explore methods for achieving objectives, including the following:**

- o *Creating new models for international talent exchanges.* International talent exchanges are powerful tools to further AI alignment, cross-pollinate ideas, and build AI-related skills and capabilities. In developing new approaches to talent exchanges, the ETC should consider:
 - Military officer exchanges to improve AI deployment and interoperability, including among NATO, JAIC, DoD, and foreign defense ministries and militaries;
 - Analogous training and exchanges needed for U.S. and allied diplomats and development experts;
 - Government-to-government exchanges of AI experts to assist in building tech and ethical expertise; exchanges to benefit industry-led multilateral and multi-stakeholder efforts like SDOs, GPAI, OECD, and influence paths taken by partners;
 - Talent exchanges and secondments in industry and academia (both international industry-industry or academia-academia talent exchanges, as well as government-industry/academia); and
 - Leveraging research centers such as the proposed MAIRI to enable cross-border collaboration and talent exchanges.
- o *Coordinating AI training development programs and sharing of best practices* for government training and broader AI education programs (including in secondary schools and universities to include computer science teaching and curriculum development).
 - The ETC should explore methods for non-EU partner nations to coordinate on the “Artificial Intelligence and Analytics” in the EU’s Digital Education Plan.

Critical Area #7 – International Digital Democracy Initiative

- The Chapter 15 Blueprint for Action details the Commission’s recommendations for coordinating foreign assistance, investment, and financing through the International Digital Democracy Initiative.

Chapter 15: A Favorable International Technology Order

Annex: Emerging Technology Coalition - Endnotes

¹ See *Interim Report and Third Quarter Recommendations*, NSCAI at 213 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

² See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendations for upholding democratic values, see the section on "Aligning Systems and Uses with American Values and the Rule of Law" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

³ Key technical standards organizations include the International Organization for Standardization (ISO), International Electrotechnical Commission (IEC), Institute of Electrical and Electronic Engineers (IEEE), the UN International Telecommunication Union's Telecommunication Standardization Sector (ITU-T), and the Third Generation Partnership Project (3GPP). See *Interim Report and Third Quarter Recommendations*, NSCAI at 205 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

⁴ See *Second Quarter Recommendations*, NSCAI at 89 (July 2020), <https://www.nscai.gov/previous-reports/>.

⁵ See *Interim Report and Third Quarter Recommendations*, NSCAI at 213 (Oct. 2020), <https://www.nscai.gov/previous-reports/>. These recommendations have been reprised and built upon in the Chapter 15 Blueprint for Action.

⁶ *Id.*

⁷ The Commission recommends that NIST provide a set of standards, performance metrics, and tools for qualified confidence in AI models, data, and training environments and predicted outcomes. The Blueprint for Action also recommends that NIST provide guidance as the science on testing across responsible AI attributes evolves. See the Chapter 7 Blueprint for Action.

⁸ Jenni Tennison, *An Introduction to the Global Partnership on AI's Work on Data Governance*, OECD AI Policy Observatory (Aug. 21, 2020), <https://oecd.ai/wonk/an-introduction-to-the-global-partnership-on-ais-work-on-data-governance>.

⁹ The OECD has led the international community with its work around AI norms and policy development. The May 2019 Principles on Artificial Intelligence was the first multilateral set of principles adopted by governments. Launched in February 2020, the OECD AI Policy Observatory facilitates dialogue between its global multi-stakeholders, provides evidence-based analysis on 20+ policy areas, promotes the adoption of the AI Principles, and bolsters the advancement and monitoring of trustworthy AI systems that benefit society. The Network of Experts on AI (ONE AI) is an informal advisory group of multidisciplinary and multi-stakeholder experts from more than 30 countries that provides policy, technical, and business expert input to inform OECD analysis and recommendations. The OECD has also developed Guidelines on the Protection of Privacy and Transborder Flows of Personal Data and is working toward principles for trusted government access to data. See generally *OECD AI Policy Observatory*, OECD.AI (last accessed Jan. 5, 2021), <https://oecd.ai/>; *OECD Privacy Guidelines*, OECD (last accessed Jan. 4, 2020), <http://www.oecd.org/digital/ieconomy/privacy-guidelines.htm>.

¹⁰ See the Technical Glossary of AI Terms Appendix of this report for definition of homomorphic encryption and differential privacy techniques. Collaborative research in this area could draw from promising R&D use cases, including the DARPA Brandeis program and the IARPA HECTOR program. See *Brandeis*, DARPA (last accessed Sept. 18, 2020), <https://www.darpa.mil/program/brandeis>; *Homomorphic Encryption Computing Techniques with Overhead Reduction (HECTOR)*, IARPA (last accessed Sept. 18, 2020), <https://www.iarpa.gov/index.php/research-programs/hector>; see the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for future R&D needed to advance capabilities for preserving and ensuring American values and the rule of law, see the section on "Aligning Systems and Uses with American Values and the Rule of Law" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹¹ See *First Quarter Recommendations*, NSCAI at 11 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹² See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for future R&D needed to advance capabilities for TEVV, see the section on "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission). TEVV R&D includes complex system testing to improve understanding of and confidence in emergent performance of composed AI systems and improve methods to understand, predict, and control systems-of-systems to avoid negative outcomes resulting from system interaction. In addition, R&D in a multi-agent scenario will advance the understanding of interacting AI systems, including the application of game theory to varied and complex scenarios and interactions between cohorts composed of a mixture of humans and AI technologies. See also *First Quarter Recommendations*, NSCAI at 11 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹³ The Commission has previously recommended that the United States devote greater resources to AI modeling, simulation, and design. See *First Quarter Recommendations*, NSCAI at 6-13 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹⁴ See the Technical Glossary of AI Terms Appendix of this report for definition of one-shot (or few-shot) learning.

¹⁵ See *First Quarter Recommendations*, NSCAI at 11 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹⁶ See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendation for future R&D needed to advance capabilities for AI security and robustness, see the section on "Engineering Practices" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission). See also *First Quarter Recommendations*, NSCAI at 11 (March 2020), <https://www.nscai.gov/previous-reports/>.

¹⁷ See the Appendix of this report containing the abridged version of NSCAI's *Key Considerations for Responsible Development & Fielding of AI*. For additional details on the Commission's recommendations to test machine-machine/multi-agent interaction and for international collaboration and cooperation needed to align on how to test and verify AI system reliability and performance along shared values, see the section on "System Performance" in *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

¹⁸ See Chapter 11 of this report and its associated Blueprint for Action for detail on potential priority areas for AI research.

¹⁹ See *Working Party on Nuclear Criticality Safety (WPNCs)*, OECD Nuclear Energy Agency, <https://www.oecd-neo.org/science/wpncs/>; *International Co-operation in Nuclear Data Evaluation: An Extended Summary of the Collaborative International Evaluated Library Organisation (CIELO) Pilot Project*, NEA No. 7489, OECD Nuclear Energy Agency (2019), <https://www.oecd-neo.org/science/wpec/documents/volume40.pdf>.

²⁰ GEOTHERMICA combines financial resources and expertise on geothermal energy research and innovation from 16 countries and their regions. It "launches joint projects that demonstrate and validate novel concepts of geothermal energy deployment within the energy system, and that identify paths to commercial large-scale implementation." One of the three focus areas includes "smart integration into the energy system and operations." Some of the projects have big-data and smart-system aspects, such as the French National Project through the GEOTHERMICA HEATSTORE project. GEOTHERMICA partners, like the U.S. Lawrence Livermore National Laboratory, have expertise in ML. See *About GEOTHERMICA*, GEOTHERMICA (last accessed Sept. 18, 2020), <http://www.geothermica.eu/about-geothermica/>; *French National Project*, HEATSTORE (last accessed Sept. 18, 2020), <https://www.heatstore.eu/national-project-france.html>; *American Partners*, GEOTHERMICA (last accessed Sept. 18, 2020), <http://www.geothermica.eu/matchmaking/united-states/>.

Chapter 15: A Favorable International Technology Order

Annex: Emerging Technology Coalition - Endnotes

²¹ Members of the partnership include the United States as well as Australia, China, Germany, Japan, Russia, Austria, Costa Rica, Iceland, ROK, South Africa, Brazil, India, the Netherlands, Canada, France, Italy, and Norway. See *International Partnership for Hydrogen and Fuel Cells in the Economy*, U.S. Department of Energy (last accessed Sept. 18, 2020), <https://www.energy.gov/eere/fuelcells/international-partnership-hydrogen-and-fuel-cells-economy>; *International Partnership for Hydrogen and Fuel Cells in the Economy*, IPHE (last accessed Sept. 18, 2020), <https://www.iphe.net/>.

²² See *Interim Report and Third Quarter Recommendations*, NSCAI at 153 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²³ Chapter 16 and its associated Blueprint for Action recommends that the United States pursue global cooperation on smart disease monitoring.

²⁴ Jason Matheny, et al., *The Role of AI Technology in Pandemic Response and Preparedness: Recommended Investments and Initiatives*, NSCAI (June 25, 2020), <https://www.nscai.gov/white-papers/covid-19-white-papers/>.

²⁵ See Chapter 11 of this report and its associated Blueprint for Action for its recommendation for the United States to take on some of humanity's biggest challenges.

²⁶ *NanoEHS CORs*, US-EU Nanotechnology Communities of Research (CORs) (last accessed Sept. 18, 2020), <https://us-eu.org/communities-of-research/>.

²⁷ See, e.g., Jeff Mason, et al., *An Overview of Clinical Applications of Artificial Intelligence*, Canadian Agency for Drugs and Technologies in Health (Sept. 2018), https://www.cadth.ca/sites/default/files/pdf/eh0070_overview_clinical_applications_of_AI.pdf.

²⁸ Feed The Future (last accessed Feb. 3, 2021), <https://www.feedthefuture.gov/>.

²⁹ *Automation and Artificial Intelligence in Agriculture: The Future of Maintaining Food Security and Sustainable Intensification*, *Frontiers* (last accessed Feb. 3, 2021), <https://www.frontiersin.org/research-topics/15206/automation-and-artificial-intelligence-in-agriculture-the-future-of-maintaining-food-security-and-su>.

³⁰ The World Economic Forum has noted 160 million people a year are at risk from natural disasters and sees great benefit in AI from "reducing the time to assess damage to monitoring social media to more quickly and effectively deliver aid" while "sharpen[ing] the decisions of relief workers on the front lines." Ashley van Heteren, et al., *Natural Disasters are Increasing in Frequency and Ferocity. Here's How AI Can Come to the Rescue*, World Economic Forum (Jan. 14, 2020), <https://www.weforum.org/agenda/2020/01/natural-disasters-resilience-relief-artificial-intelligence-ai-mckinsey/>.

³¹ David Vergun, *DOD Partners With Agencies to Use AI for Disaster Relief, Humanitarian Relief*, DOD News (Aug. 20, 2020), <https://www.defense.gov/Explore/News/Article/Article/2319945/dod-partners-with-agencies-to-use-ai-for-disaster-humanitarian-relief/>.

³² Prashanth Parameswaran, *What's in the New US-Singapore Artificial Intelligence Defense Partnership?*, *The Diplomat* (July 1, 2019), <https://thediplomat.com/2019/07/whats-in-the-new-us-singapore-artificial-intelligence-defense-partnership/>.

³³ See the discussion in Chapter 11 and its associated Blueprint for Action on using AI to tackle some of humanity's biggest challenges.

³⁴ On asteroids, see *Deep Asteroid*, NASA (May 27, 2016), <https://open.nasa.gov/innovation-space/deep-asteroid/>.

³⁵ *AI Challenged to Stave off Collisions in Space*, European Space Agency (Oct. 9, 2019), https://www.esa.int/Enabling_Support/Space_Engineering_Technology/AI_challenged_to_stave_off_collisions_in_space; Angelica Mari, *UK Government Seeks Innovations to Tackle Space Debris*, *Computer Weekly* (May 28, 2020), <https://www.computerweekly.com/news/252483762/UK-government-seeks-innovations-to-tackle-space-debris>.

³⁶ Andrew Myers, *Stanford Develops an AI Navigation System for a Future Satellite 'Tow Truck,'* Stanford News (Feb. 1, 2019), <https://news.stanford.edu/2019/02/01/stanford-spurs-ai-navigation-space-rendezvous-software/>.

³⁷ Leslie D'Monte, *Chandrayaan-2 Pragyan Shows How AI is Helping Space Exploration*, Mint (Sept. 6, 2019), <https://www.livemint.com/technology/tech-news/chandrayaan-2-pragyan-shows-how-ai-is-helping-space-exploration-1567764065716.html>.

³⁸ Mike Wall, *New, Emotionally Intelligent Robot CIMON 2 Heads to Space Station*, Space.com (Dec. 5, 2019), <https://www.space.com/cimon-2-artificial-intelligence-robot-space-station.html>.

³⁹ The EU intends to establish Centres of Excellence and Digital Innovation Hubs focused on AI. ICEMAI works with the Government of Canada's Advisory Council on Artificial Intelligence, Forum IA Quebec, and the International Observatory on the Societal Impacts of Artificial Intelligence and Digital Technologies and is supported by the governments of Canada and Quebec with up to \$15 million in funding over five years. INRIA was launched in February 2020 and has a contract with the Government of France to focus on "speeding up development of France's scientific and technological leadership, as part of a Europe-wide approach," including prioritizing AI and other digital technologies to meet societal challenges, constructing European research and innovation spaces, strengthening the tech industrial base, reinforcing public policies, and developing leading research universities. See *Communication Artificial Intelligence for Europe*, European Commission (April 25, 2018), <https://ec.europa.eu/digital-single-market/en/news/communication-artificial-intelligence-europe>; *The Global Partnership on Artificial Intelligence Officially Launched*, Montreal International (June 15, 2020), <https://www.montrealinternational.com/en/news/the-global-partnership-on-artificial-intelligence-officially-launched/>; *INRIA: For Scientific, Technological and Industrial Leadership in Digital Technology*, Government of France (Feb. 24, 2020), <https://www.gouvernement.fr/en/inria-for-scientific-technological-and-industrial-leadership-in-digital-technology>.

⁴⁰ See Implementation Plan to Chapter 15: A Favorable International Digital Order.

⁴¹ *AI To Solve Global Issues*, XPRIZE (last accessed Sept. 18, 2020), <https://www.xprize.org/prizes/artificial-intelligence>.

⁴² Melissa Flagg & Paul Harris, *System Re-engineering: A New Policy Framework for the American R&D System in a Changed World*, Center for Security and Emerging Technology (Sept. 2020), <https://cset.georgetown.edu/research/system-re-engineering/>.

⁴³ The National AI Research Institutes is a joint government effort among the National Science Foundation (NSF), U.S. Department of Agriculture (USDA) National Institute of Food and Agriculture (NIFA), U.S. Department of Homeland Security (DHS) Science & Technology Directorate (S&T), and the U.S. Department of Transportation (DOT) Federal Highway Administration (FHWA). See *National Artificial Intelligence (AI) Research Institutes*, NSF (last accessed Sept. 18, 2020), https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=505686.

⁴⁴ See Chapter 8 of this report and its associated Blueprint for Action for its recommendation for Congress to authorize NIST to sponsor a University Affiliated Research Center (UARC), Federally Funded Research & Development Center (FFRDC), and/or a lab to provide independent, third-party testing.

⁴⁵ The AI4EU project was founded by the European Commission under the H2020 program to establish the first European AI On-Demand Platform and Ecosystem. The Platform is designed to support the AI ecosystem and provide a forum to share AI resources from European projects. See *About the Project*, AI4EU (last accessed Sept. 18, 2020), <https://www.ai4eu.eu/about-project>.

⁴⁶ *AI for Humanity*, Mila (last accessed Sept. 18, 2020), <https://mila.quebec/ia-dans-la-societe/>.

⁴⁷ *Artificial Intelligence For Good and For All*, Amii (last accessed Sept. 18, 2020), <https://www.amii.ca/>.

⁴⁸ See *Pan-Canadian AI Strategy*, CIFAR (last accessed Feb. 4, 2021), <https://www.cifar.ca/ai>; *About Us*, Vector Institute for Artificial Intelligence (last accessed Sept. 18, 2020), <https://vectorinstitute.ai/about/>.

Chapter 15: A Favorable International Technology Order

Annex: Emerging Technology Coalition - Endnotes

⁴⁹ See Stanford Institute for Human-Centered Artificial Intelligence (last accessed Sept. 18, 2020), <https://hai.stanford.edu/welcome>.

⁵⁰ The U.S. has strong research ties to Israel through the Binational Science Foundation (BSF) and the Binational Industrial Research & Development Foundation (BIRD). See *About the BSF*, U.S.-Israel Binational Science Foundation (last accessed Feb. 3, 2021), <https://www.bsf.org.il/about/>; *What is BIRD?* U.S.-Israel Binational Industrial Research and Development (last accessed Feb. 3, 2021), <https://www.birdf.com/what-is-bird/>. The Indo-U.S. Science and Technology Forum (IUSSTF) oversees the United States-India Science & Technology Endowment Fund (USISTEF), which supports and fosters joint applied R&D. *About the Fund*, Indo-U.S. Science and Technology Forum (last accessed Feb. 3, 2021), <https://www.iusstf.org/usistef/us-india-science-technology>.

⁵¹ See Joint Statement on Artificial Intelligence and Human Rights, Freedom Online Coalition (last accessed Jan. 5, 2021), <https://freedomonlinecoalition.com/wp-content/uploads/2020/11/FOC-Joint-Statement-on-Artificial-Intelligence-and-Human-Rights.pdf>.

⁵² The FY 2021 NDAA created the Open Technology Fund as Section 309A of the U.S. International Broadcasting Act of 1994. Pub. L. 116-283, sec. 1299P, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021). Congress has since appropriated \$20 million to the fund. See Joint Explanatory Statement, Division K— Department of State, Foreign Operations, and Related Programs Appropriations Act, 2021 at 23 (2021), <https://www.appropriations.senate.gov/imo/media/doc/Division%20K%20-%20SFOPS%20Statement%20FY21.pdf> (enacted in Pub. L. 116-260, the Consolidated Appropriations Act, 2021). The Department of State's Internet Freedom and Business & Human Rights Section (IFBHR), within DRL, leads U.S. Government policy and engagement to protect human rights online. See *Internet Freedom: Fact Sheet*, U.S. Department of State (Nov. 17, 2017), <https://2017-2021.state.gov/internet-freedom/index.html>. IFBHR works across the U.S. Government, with democratic nations, with civil society, and with the Freedom Online Coalition. IFBHR's program includes funding development of censorship-defeating peer-to-peer communications technologies. See *Internet Freedom: Advancing and Promoting Peer-to-Peer Communications Technologies*, U.S. Department of State (Feb. 13, 2020), <https://2017-2021.state.gov/internet-freedom-advancing-and-promoting-peer-to-peer-communications-technologies/index.html>.

⁵³ Isaac Ben-Israel, et al., *Towards Regulation of AI Systems: Global Perspectives on the Development of a Legal Framework on Artificial Intelligence (AI) Systems Based on the Council of Europe's Standards on Human Rights, Democracy and the Rule of Law*, CAHAI Secretariat (Dec. 2020), <https://rm.coe.int/prems-107320-gbr-2018-compli-cahai-couv-texte-a4-bat-web/1680a0c17a>.

⁵⁴ See the Chapter 14 Blueprint for Action. See also *Guidance on Implementing the UN Guiding Principles for Transactions Linked to Foreign Government End-Users for Products or Services with Surveillance Capabilities*, U.S. Department of State (last accessed Jan. 4, 2021), <https://www.state.gov/wp-content/uploads/2020/10/DRL-Industry-Guidance-Project-FINAL-1-pager-508-1.pdf>.

⁵⁵ See the Chapter 1 Blueprint for Action on malign information for further details on this proposal.

⁵⁶ *About*, Global Internet Forum to Combat Terrorism (last accessed Jan. 5, 2021), <https://www.gifct.org/about/>.

⁵⁷ Civil society in Taiwan has responded to the threat from disinformation in a number of ways, including demonstrating outside compromised media firms, educating senior citizens on the ways they may be exposed to disinformation, and the establishment of robust fact-checking groups such as the Taiwan Fact Check Center, MyGoPen, Cofacts, and Rum Toast. These groups have worked with both government and social media platforms to not only identify and remove disinformation, but also to forensically trace disinformation back to sources in China. See *Audrey Tang on Taiwan's Digital Democracy, COVID-19, and Combating Disinformation*, The Stimson Center (March 18, 2020), <https://www.stimson.org/2020/interview-with-taiwan-digital-minister-audrey-tang/>.

⁵⁸ For more on democratic use of surveillance technologies, see Chapter 8 of this report, Upholding Democratic Values: Privacy, Civil Liberties, and Civil Rights in Uses of AI for National Security.

⁵⁹ Kara Frederick, *The Razor's Edge: Liberalizing the Digital Surveillance Ecosystem*, CNAS (Sept. 3, 2020), <https://www.cnas.org/publications/reports/the-razors-edge-liberalizing-the-digital-surveillance-ecosystem>.

⁶⁰ See *OpenMined's Efforts for the Coronavirus Pandemic: COVID Alert App, Private Set Interaction, A Differential Privacy Wrapper and Private Identity*, OpenMined (April 1, 2020), <https://blog.openmined.org/openmineds-efforts-for-the-coronavirus-pandemic/>.

⁶¹ The U.S. and Europe should agree on a common definition for anonymized data to include a clearer understanding of what constitutes “consent” and “matters of public interest.”

⁶² Remarks by Angel Gurría, OECD Secretary General, delivered at the 2019 G20 Leaders' Summit—Digital (AI, data governance, digital trade, taxation) (June 28, 2019), <https://www.oecd.org/g20/summits/osaka/2019-g20-leaders-summit-digital-osaka-june-2019.htm>.

⁶³ *A New EU-US Agenda for Global Change*, European Commission (Dec. 2, 2020), https://ec.europa.eu/info/sites/info/files/joint-communication-eu-us-agenda_en.pdf.

⁶⁴ The Commission's *Key Considerations* and existing international principles could be leveraged, such as the OECD Principles on AI, the OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, the OECD Recommendation on Digital Security of Critical Activities, the forthcoming OECD Principles on Trusted Government Access to Data, and others.

⁶⁵ See Chapter 14 of this report and its associated Blueprint for Action for additional details on the Commission's recommendations regarding aligning allied export control and investment screening regimes. Within the U.S. government, the Departments of State and Commerce, on export controls, and the Departments of State and the Treasury, on investment screening, have already begun such work.

⁶⁶ See the Chapter 15 Blueprint for Action.

⁶⁷ See *Second Quarter Recommendations*, NSCAI at 63-67 (July 2020), <https://www.nscai.gov/previous-reports/>. In particular, the United States, the Netherlands, and Japan should coordinate export controls on extreme ultraviolet and ArF immersion lithography tools, as doing so would limit the ability of China and other competitors to develop the high-end microelectronics that are increasingly essential for AI. For additional details on the Commission's recommendations regarding export controls on SME, see Chapter 14 of this report.

⁶⁸ See Chapter 13 of this report for additional details and recommendations on the microelectronics supply chain.

⁶⁹ See Chapters 3 and 14 of this report.

⁷⁰ See *Second Quarter Recommendations*, NSCAI at 48 (July 2020), <https://www.nscai.gov/previous-reports/>; Andrew Imbrie et al., *Agile Alliances: How the United States and Its Allies Can Deliver a Democratic Way of AI*, Center for Security and Emerging Technology at 16-17 (Feb. 2020), <https://cset.georgetown.edu/research/agile-alliances/>; Andrew Imbrie, et al., *The Question of Comparative Advantage in Artificial Intelligence: Enduring Strengths and Emerging Challenges for the United States*, Center for Security and Emerging Technology at 33 (Jan. 2020), <https://cset.georgetown.edu/wp-content/uploads/CSET-The-Question-of-Comparative-Advantage-in-Artificial-Intelligence-1.pdf>.

⁷¹ *First Quarter Recommendations*, NSCAI at 45 (March 2020), <https://www.nscai.gov/previous-reports/>.

⁷² See Chapters 13 and 16 of this report.

⁷³ “IP5” is the name of the forum of the world's five largest IP offices that was set up to improve the efficiency of the examination process for patents worldwide. See *About IP5 Co-Operation*, fiveIPoffices (last accessed Jan. 4, 2021), <https://www.fiveipoffices.org/about>.

⁷⁴ See Chapter 12 of this report and its associated Blueprint for Action.

Chapter 15: A Favorable International Technology Order

Annex: Emerging Technology Coalition - Endnotes

⁷⁵ See Chapter 12 of this report and its associated Blueprint for Action.

⁷⁶ See Chapter 10 of this report.

⁷⁷ This approach has been recommended by Melissa Flagg and Zachary Arnold. See Melissa Flagg & Zachary Arnold, *A New Institutional Approach to Research Security in the United States Defending a Diverse R&D Ecosystem*, Center for Security and Emerging Technology (Jan. 2021), <https://cset.georgetown.edu/research/a-new-institutional-approach-to-research-security-in-the-united-states/>.

⁷⁸ See Chapter 1 and its associated Blueprint for Action on preparing for AI-enabled cyber conflict.

Chapter 16:

Associated Technologies

Blueprint for Action

Recognizing that leadership in artificial intelligence (AI) relies on leadership across a suite of emerging technologies, the United States must prioritize the research and development (R&D), application, and adoption of not just AI, but the technologies that enable it and are enabled by it. This process should be based on a careful analysis of the national security threats and opportunities at the intersection of AI and its associated technologies. If the U.S. government fails to adopt a more strategic approach to protecting and promoting U.S. advantages in these areas, it risks jeopardizing the country's technological leadership, economic prosperity, and national security.

In accordance with its mandate to consider both AI and “associated technologies,” the Commission identifies and proposes steps to maintain U.S. leadership across the spectrum of technologies it believes are most critical to U.S. national competitiveness. The Commission then offers specific recommendations on how the United States can proactively address the novel national security threats and opportunities posed by three technologies in particular: biotechnology, quantum computing, and 5G telecommunications.¹ Finally, the Commission expands its analysis to include recommendations on a broader set of emerging technologies critical to U.S. national competitiveness.

Technologies Critical to U.S. National Competitiveness

The Commission has identified eight technologies and related platforms that are key to U.S. leadership. Two of these technologies—AI and microelectronics—are addressed elsewhere in this report. The remaining six—biotechnology, quantum computing, 5G and advanced networking, autonomy and robotics, advanced and additive manufacturing, and energy systems—are covered below. These recommendations build on the Commission's previous work by providing actions the U.S. government could take to promote overall U.S. leadership and long-term competitiveness across the constellation of emerging technologies.

Recommendation

Recommendation: Identify and Prioritize Technologies Central to National Competitiveness

To date, there is no whole-of-government consensus for which emerging technologies are most critical to long-term strategic competitiveness and whose development must

be prioritized. Several government agencies have made independent attempts to define such a list: the 2018 National Defense Strategy,² the list of “critical emerging technologies” produced by the Department of Defense in response to Section 1793 of the FY 2019 National Defense Authorization Act (NDAA),³ the Department of Commerce’s 2018 Advance Notice of Proposed Rulemaking (ANPRM) of controls on certain emerging technologies,⁴ the report by the President’s Council of Advisors on Science and Technology titled *Recommendations for Strengthening American Leadership in Industries of the Future* from 2020,⁵ and the bill introduced by Senator Chuck Schumer in the 116th Congress with seven bipartisan co-sponsors titled the “Endless Frontier Act.”⁶ Additionally, the White House published the *National Strategy for Critical and Emerging Technologies* in October 2020, which included a list of critical and emerging technologies.⁷ However, this document does not explain why each of these technologies is essential to U.S. national competitiveness, nor does it include specific implementation plans for promoting their development and protecting U.S. advantages in each.

These lists have substantial overlap, but no two lists are the same and no single list is authoritative. Consequently, there is no whole-of-government consensus, and certainly no national consensus, of which technologies are critical to U.S. national competitiveness, making it more difficult for the U.S. government to marshal private-sector investment, for legislators to prioritize funding, and for U.S. government agencies to coordinate technology protection and promotion. There is also no list around which the White House can organize a national technology strategy and no coordinated mechanism within the U.S. government to support financing of these priorities when there are market failures and private-sector financing is insufficient.

Actions for the Executive Office of the President:

- **Define and prioritize the key emerging technologies in which U.S. leadership is essential.**
 - o The Executive Office of the President, in consultation with departments and agencies, should publish a single, authoritative list of technologies and sectors which are key to overall U.S. competitiveness, along with detailed implementation plans for each to ensure long-term U.S. leadership.
 - The implementation plans should identify specific subcomponents of each technology that are most important, key choke points where competitors could be blocked with minimal impact on U.S. industry, and where additional resources are needed. These plans should include specific steps to promote domestic industry, ensure supply chain resiliency, and protect key technologies from competitors. This list of technologies and the associated implementation plans will form the core of a National Technology Strategy, as referenced in Chapter 9 of this report.
 - o The creation and maintenance of such a list and implementation plans will help produce a national consensus regarding which industries are most important in the emerging techno-economic competition. The result will be an important message to Congress regarding where the country must prioritize and expend resources,

as well as a powerful demand signal to industry. The figure below includes eight technologies that the Commission recommends be considered for the list.

U.S. Government Lists of Critical Technologies						
NSCAI-Proposed Critical Technology List	2018 National Defense Strategy	DoD List of Critical Emerging Technologies	Commerce ANPRM on Emerging Technologies	PCAST List of Industries of the Future	S.3832 - Endless Frontier Act	WH Nat Strategy for C&ET
Artificial Intelligence	✓	✓	✓	✓	✓	✓
Biotechnology	✓	✓	✓	✓	✓	✓
Quantum Computing		✓	✓	✓	✓	✓
Semiconductors and Advanced Hardware	✓	✓	✓		✓	✓
Autonomy and Robotics	✓	✓	✓		✓	✓
5G and Advanced Networking		✓		✓	✓	✓
Advanced Manufacturing			✓	✓	✓	✓
Energy Systems	✓	✓			✓	✓

- **Expand the loan authority of the Development Finance Corporation to include domestic industrial base capabilities supporting key emerging technologies.**
 - o The President should issue an Executive Order that expands the loan authority of the U.S. International Development Finance Corporation (DFC) to include domestic industrial base capabilities related to any of the aforementioned technologies that are identified by the Executive Office of the President as key to overall U.S. competitiveness.
 - Specifically, the Executive Order should delegate authority under Title III of the Defense Production Act to the DFC to issue loans that “create, maintain, protect, expand, or restore domestic industrial base capabilities” supporting the aforementioned list of technologies, or “the resiliency of relevant domestic supply chains.” This new authority should be of indefinite duration.
 - This action would build off of Executive Order 13922, which expanded similar domestic loan authorities to DFC related to industries supporting “the national response and recovery to the COVID-19 outbreak” until 2022.⁸

- o Expanding the domestic authorities of the DFC as it relates to critical technologies will help the government support key platforms and projects which are critical to future U.S. national security and economic competitiveness but lack sufficient private-sector capital.
 - The DFC should coordinate with the Technology Competitiveness Council recommended in Chapter 9 of this report to identify specific platforms that are most in need of such financing.

Ensuring U.S. Leadership in Biotechnology

The combination of advances in AI and biology have the potential to reshape the global economy for the next century. Progress in genetic sequencing has given researchers the ability to read the “code of life.” Given the significant quantity of data involved, AI will be essential to fully understanding how genetic code interacts with biological processes. Finally, advances in synthetic biology and genetic editing will give researchers the ability to manipulate this code to perform specific functions. Together, these techniques will enable transformational breakthroughs in biology and underpin most future scientific breakthroughs related to human health, agriculture, and climate science. The nation which is best able to simultaneously leverage both technologies will have substantial strategic advantages for the foreseeable future, potentially becoming a global leader in pharmaceuticals, reducing its reliance on foreign supply chains, and even ensuring it has a healthier and more capable population. These technological breakthroughs will also cause the biotechnology sector to become a major driver of overall U.S. economic competitiveness.

Recommendation: Prioritize the Development of an Advanced Biotechnology R&D Ecosystem

Recommendation

The United States must invest in key platforms that better position the U.S. academic and commercial biotech industry to benefit from AI-enabled advancements in biology. It should specifically look to support platforms that aggregate biodata, and specifically genetic data, in a secure manner in order to enhance the ability of U.S. researchers to utilize AI to facilitate breakthrough biotechnology research and innovation. Additionally, the United States should support efforts to expand the scope and sophistication of U.S. biofabrication capabilities to ensure it can keep pace with forthcoming research advancements. It should specifically support efforts to transform the biotechnology industry away from its current, vertically integrated models and encourage the development of multiple standardized, merchant biofabrication facilities. Doing so would expand access to advanced biofabrication tools among startups and laboratories by allowing firms to rapidly design new molecules and materials via the cloud and place immediate orders for fabrication.

Actions for Congress and the Department of Health and Human Services:

- **Fund and establish a world-class biobank for genetic data.**
 - o Congress should fund efforts to build a world-class biobank within the National Institutes of Health (NIH). The current leading U.S. genetic database, GenBank, is underfunded, difficult to access, and poorly curated, particularly in comparison to other leading genetic databases such as the U.K. BioBank or the China National GeneBank. The entity should be securely and easily accessible by legitimate researchers; contain a wide variety of whole human, animal, and plant genomes, including de-identified metadata about phenotypes; and aggregate other open and potentially even proprietary datasets for specialized uses. It must also include strong privacy protections for human genetic data. Creating and staffing such an entity would likely require a budget of approximately \$100 million per year, on top of up-front construction costs.⁹
- **Direct funding to support advanced biotech manufacturing initiatives through entities such as BARDA.**
 - o The Department of Health and Human Services should direct funds to support advanced biotech manufacturing initiatives through entities such as the Biomedical Advanced Research and Development Authority (BARDA), and Congress should prioritize such initiatives in future health-related spending bills. This could take the form of financial incentives for advanced biotech manufacturing firms focused on sophisticated, flexible, cloud-based fabrication, or R&D funding to support advanced manufacturing techniques.

Recommendation

Recommendation: Prioritize Advanced Biotechnology Capabilities as Imperative for National Security and Economic Competitiveness

The growing importance of biotechnology leadership to health, food, production, and science also makes it a national security imperative that the United States take proactive steps to facilitate long-term U.S. leadership in the field. Advancements in biotechnology will also create novel national security challenges, ranging from engineered pathogens to augmented competitor human physiological or mental capabilities. The United States currently is not postured to address such challenges, and biological threats have rarely been a priority issue for the U.S. national security community. The COVID-19 pandemic clearly illustrates that the United States must think more broadly about national security threats than it has in the past, and that biological threats in particular have the potential to impose significant costs on U.S. society and security.

U.S. competitors see the potential for AI to spur new, transformational advances in biotechnology. China in particular is actively seeking global leadership in both fields, sees its AI and biotechnology strategies as mutually reinforcing, and believes the synergies between the two will translate into military advantage.¹⁰ China also faces fewer barriers to collecting, using, and combining human biological data given its disregard for individual privacy and bioethical principles. The global reach of China's genomics giant, BGI, poses similar threats in the biotechnology sector as Huawei does in the communications sector.

Actions for the Executive Office of the President:

- **Update the U.S. National Biodefense Strategy to include additional AI-enabled biological threats.**¹¹
 - The National Security Council should update its *National Biodefense Strategy*, which currently only focuses on natural or engineered pathogens, to include a wider vision of biological threats.¹² The strategy should specifically examine how AI could enable new biological advances which pose unique national security threats, such as human enhancement, and how U.S. competitors could utilize advantages in biotechnology or biodata as an instrument of national power. It should also specifically consider how AI could identify and counter the creation of advanced, engineered pathogens which target certain elements of the U.S. population or food supply. AI is facilitating a rapid evolution of the biotechnology field, and the U.S. biodefense strategy must evolve with it.
- **Direct departments and agencies to prioritize initiatives that promote U.S. biotechnology leadership.**
 - Directing departments and agencies to prioritize initiatives promoting U.S. biotechnology leadership would include aggressively promoting funding for basic research in biology, particularly applications of biology that utilize AI; focusing resources on forecasting how AI will enable future biotechnology breakthroughs; and continuing to cultivate talent both inside and outside the government, as well as commercial activity at the nexus of AI and biology. This will require an entity which is empowered to coordinate across the economic, technological, and security spheres, such as the Commission's recommended Technology Competitiveness Council.¹³

Recommendation: Publicly Highlight BGI's Links to the Chinese Government

Recommendation

BGI is China's de facto national champion in genetic sequencing and research and is among the world leaders in DNA sequencing. It has research affiliations with multiple U.S. universities, including the University of Washington and Washington State University.¹⁴ BGI has also benefited from substantial support from the Chinese government, as well as its 2013 acquisition of a competing U.S. firm, Complete Genomics.¹⁵ There are indications that BGI's links with the Chinese government may run deeper than it publicly claims, as it built and operates China National GeneBank, the Chinese government's national genetic database, and has used PLA-owned supercomputers to process genetic information.¹⁶ Chinese diplomats have pushed BGI-built COVID-19 testing kits, including in the United States, and by August 2020 BGI had "sold 35 million rapid COVID-19 testing kits to 180 countries, and built 58 labs in 18 countries."¹⁷

BGI may be serving, wittingly or unwittingly, as a global collection mechanism for Chinese government genetic databases, providing China with greater raw numbers and diversity of human genome samples as well as access to sensitive personal information about key individuals around the world. The highest levels of the United States government should publicly state these concerns so as to raise awareness among the U.S. commercial and

academic biotechnology communities, as well as U.S. allies, many of which currently have partnerships or business dealings with BGI.

Action for the Department of State:

- **Launch a strategic communications campaign to publicly highlight the links between the Chinese government and BGI.**

- o The Secretary of State should personally voice concern about BGI's ties to the Chinese government and instruct the Department to conduct a strategic communications campaign to highlight those links and warn of the dangers of the Chinese government obtaining personal genetic information via BGI. The Department should also warn BGI and the Chinese government that it will closely monitor BGI's activities, and that should BGI be utilized as a mass DNA-collection apparatus for the Chinese government it could face additional U.S. regulatory action.

Recommendation

Recommendation: Pursue Global Cooperation on Smart Disease Monitoring

While pivoting to a more competitive national approach toward biotechnology policy, the United States should also pursue efforts to enhance global cooperation on disease monitoring. By pooling existing open-source health-related data with improved early warning signals and data on zoonotic spillovers and transmission of novel viruses, governments will be better postured to use AI to predict and contain future pandemics. Combining increased transparency and data sharing on disease outbreaks with AI tools—which can enhance early outbreak detection and contribute to real-time disease monitoring—could provide substantial benefit for global public health if all countries, including China, participated in good faith.¹⁸

Action for the Departments of State and Health and Human Services:

- **Support multilateral efforts to promote smart disease monitoring.**

- o The Departments of State and Health and Human Services should lead and support multilateral efforts to promote smart disease monitoring. In particular, the United States should pursue efforts to integrate and standardize international health-related data sets and combine them with global data about zoonotic spillovers to allow for the utilization of AI technologies to create shared, predictive, global disease-monitoring tools.

Ensuring U.S. Leadership in Quantum Computing

Quantum computing has the potential to create new national security threats and opportunities by enhancing the speed and precision of existing AI systems and creating new capabilities that could fundamentally alter the strategic environment. For example, quantum computers may be able to more efficiently optimize logistics for the military or

discover new materials for weapon systems.¹⁹ Quantum sensors and communications are also poised to revolutionize the collection and transfer of sensitive information, which directly affects how AI is trained and deployed in national security use cases.²⁰ Failure to step up investment in the R&D of materials and components for quantum computers, open-source software tools, and hybrid quantum-classical algorithms that leverage noisy intermediate-scale quantum computers may leave the United States vulnerable to strategic surprise on behalf of competitors.²¹

Recommendation: Transition from Basic Research to National Security Applications of Quantum Computing

Recommendation

Although the United States is well-positioned to take advantage of its early success in the basic science of quantum computing, the U.S. Government must increase its focus on fielding national security applications or risk falling behind strategic competitors. Most notably, China has made significant investments in military applications of quantum computing in an attempt to offset U.S. strengths.²² The Department of Defense (DoD) is still refining its approach to rapidly transition commercial technologies from research to fielding in high-cost, hardware-intensive sectors such as quantum computing. In the long term, DoD should prioritize efforts to rapidly procure technology across its innovation offices, but this process could take several years of dedicated effort. In the interim, announcements of priority applications will help spur private-sector investment and innovation in quantum computing despite the absence of an integrated technology-procurement apparatus.²³

Action for the President:

- **The President should direct departments and agencies to announce priority use cases of quantum computers.**
 - o The National Quantum Coordination Office (NQCO) should coordinate an effort by departments and agencies represented on the National Science and Technology Council (NSTC) Subcommittee on Quantum Information Science (QIS) to announce their priority use cases of quantum computers. By reflecting the combined views of federal entities engaging with the private sector, this would signal that a market for practical applications of quantum computing exists, set clear and specific goals for the private sector to pursue, and incentivize additional private investment. Some applications of quantum computers may be too sensitive to reveal publicly, but those that can be announced will provide direction to the private sector and facilitate the commercialization of quantum computers, which can then be applied to national security use cases.

Recommendation: Foster a Vibrant Domestic Quantum Fabrication Ecosystem

Recommendation

Due to the strategic implications of quantum computing and its application to AI, the United States must take steps now to cement its long-term status as the global leader in the design and manufacturing of quantum processing units (QPUs). To avoid the situation in which the U.S. semiconductor industry currently finds itself, the United States must

establish trusted and assured sources for critical materials and components of QPUs, ranging from manufacturing equipment to superconductors and dilution refrigerators.²⁴ Although these materials and components may not yet represent choke points, they will inevitably become more specialized as the manufacturing processes required to design and produce QPUs continue to advance. Rather than reshoring the entire supply chain for QPUs, the United States should work with its allies to develop a resilient network of suppliers for critical components that directly impact U.S. national security.

However, a secure supply chain is not sufficient to ensure U.S. leadership in quantum computing. To benefit from future breakthroughs in the field, the United States must create a robust domestic ecosystem for the research, development, and application of quantum computers that attracts top-tier talent from around the world.²⁵ The U.S. Government should offer incentives for the R&D of quantum computers and their components while simultaneously creating demand for national security applications of quantum technologies. The Quantum Economic Development Consortium (QED-C), proposed in the National Quantum Initiative (NQI) Act of 2018, is an important step toward extending U.S. leadership in next-generation computer hardware for years to come.²⁶

Action for Congress:

- **Enact a package of provisions that incentivizes the domestic design and manufacturing of quantum computers and their constituent materials.**
 - o A tax credit for expenditures made in the United States on research and development, manufacturing equipment, and workforce training related to the development of quantum computers is a necessary, albeit not sufficient, step to maintain U.S. competitiveness in this area. This provision could be modeled on the Alternative Simplified Credit (ASC), which provides a credit of 14% of expenditures on R&D in excess of 50% of base period expenditures. To help startups on the cutting edge of research and development access funding that allows them to scale, the U.S. Government should also provide loan guarantees and equity financing.

Recommendation

Recommendation: Make Quantum Computing Accessible to Researchers via the National AI Research Resource (NAIRR)

Despite recent advances in the fields of quantum hardware and software, fault-tolerant quantum computers (FTQCs) capable of performing general-purpose tasks are unlikely to replace classical computers anytime soon. In the near term, the United States should invest in noisy intermediate-scale quantum (NISQ) computers that are capable of deriving probabilistic solutions from imperfect qubits.²⁷ Hybrid quantum-classical techniques have also shown promise, whereby classical computers delegate certain tasks to purpose-built quantum devices within the same workflow. However, resources suitable for developing this type of software are not readily accessible.²⁸ By making classical and quantum computers available in the same workflow, the U.S. Government would lower barriers to innovation for

startups in the quantum computing space and attract top-tier talent from around the world. The resulting public-private partnerships would also encourage the commercialization of quantum computers and help the U.S. Government adopt those products for national security use cases.

Action for the Executive Branch:

- **Make classical and quantum computers available in the same workflow via the National AI Research Resource.**
 - By providing access to both classical and quantum computers via the National AI Research Resource (NAIRR), which the Commission recommended establishing in its *First Quarter Recommendations* and describes in greater detail in Chapter 11 of this report,²⁹ the U.S. Government would help researchers from industry, academia, and government build and test software tools and algorithms that leverage both classical and quantum computers in a hybrid fashion. These types of applications are likely to be the nearest-term use case of quantum computers.

Ensuring U.S. Leadership in 5G Telecommunications

AI systems require high-fidelity sensing as well as fast, safe, and secure networks. It is a national security imperative for the U.S. military and the nation as a whole to have access to a powerful 5G network to enable future AI capabilities and ensure the network is trusted. The United States must preserve this access and trust while building out commercial 5G networks domestically and internationally.

Recommendation: Accelerate U.S. 5G Deployment Through Spectrum Sharing

Recommendation

The slow rollout of 5G networks in the United States compared to China risks undermining U.S. advances in AI, both in the government and the private sector.³⁰ The sub-6 GHz spectrum, sometimes referred to as the mid-band or the “goldilocks” band of spectrum, is the critical portion of the spectrum for both DoD and commercial 5G operations. Sub-6 GHz spectrum is critical for 5G civilian communications since it combines high data rates with good range and penetration. Within DoD, it is also already used by many radar and communication systems because it also combines high discrimination capability with long-range operations.³¹ In part due to its importance to military operations, DoD has retained exclusive access to significant portions of the mid-band spectrum, which limits commercial uses. Unfortunately, the lack of U.S. mid-band spectrum commercial availability is substantially slowing the deployment of 5G networks domestically. Given that sub-6 GHz is important for sensing using radar and civilian communications, spectrum sharing between DoD and the private sector is the ideal approach to enabling access for both purposes in a manner that balances national security and economic interests.³²

Several U.S. Government agencies are working to address this problem by developing spectrum-sharing capabilities within the 3- to 6-GHz range. In 2015, the Federal

Communications Commission (FCC) established the Citizens Broadband Radio Service (CBRS), the first U.S. spectrum sharing model.³³ Since that time, the National Telecommunications and Information Administration (NTIA) has studied, and has collaborated with the DoD and FCC on, maximizing spectrum-sharing capabilities.³⁴ The CBRS enables shared federal and non-federal use of the band. This work allows the U.S. Navy and non-government providers to share the 3550-3700 MHz band across three dynamically managed tiers: the Navy will maintain first priority access, followed by companies and organizations that purchase priority-access licenses, and finally companies and organizations that register at no cost. The FCC held its first auction for priority-access licenses for this band in July 2020, which raised more than \$4.5 billion through the sale of 20,625 licenses.³⁵ This is a promising but modest start and these efforts must expand to a larger portion of the mid-band spectrum to be competitive with China. To achieve spectrum sharing at a competitive level will require technical analysis and engagement with industry. A comprehensive process will be critical to ensuring that DoD maintains access to spectrum essential for operational effectiveness while also broadening commercial access to spectrum for civilian 5G networks.³⁶

Action for the NTIA, FCC, and DoD:

- **Expand spectrum-sharing programs led by NTIA, FCC, and DoD, starting with a one-year 5G spectrum-sharing demonstration program.**
 - o The Commission urges NTIA, the FCC, and DoD to jointly expand spectrum-sharing programs such as the CBRS and work to license additional sub-6GHz spectrum to wireless carriers and equipment makers for commercial 5G use. Sharing and licensing additional mid-band spectrum will ensure unrestricted DoD access in the event of an emergency while also opening up 5G for commercial use. However, current spectrum-sharing capabilities must be further analyzed, tested, and demonstrated before they can be scaled. The Commission supports a one-year demonstration program that includes NTIA, FCC, DoD, and industry to assess the network's capabilities and its capacity to dynamically share spectrum between government and civilian users. If successful, such a network would be rapidly scaled with commercially available equipment.

Promote U.S. Leadership in Other Key Emerging Technologies

AI, microelectronics, biotechnology, quantum computing, and 5G telecommunications are not the only emerging technologies that will underpin U.S. national competitiveness in the 21st century. The Commission assesses that the full spectrum of emerging technologies key to U.S. technological leadership extends further and includes autonomy and robotics, advanced manufacturing, and energy systems. The Commission therefore recommends several actions to ensure U.S. leadership in these additional key emerging technologies.³⁷

Autonomy and Robotics

Recommendation: Incentivize the Development of World-Class Software Platforms for Robotic and Autonomous Systems

Recommendation

Autonomous systems that rely on robotics to execute tasks in the real world are being applied to everything from advanced manufacturing to warfighting.³⁸ As AI continues to improve the ability of these systems to match or exceed human capabilities, the United States must position itself as a leading producer and adopter of robotic hardware and software for civilian and military use cases. The United States currently lags behind countries such as Japan and Korea on the manufacturing and installation of industrial robots, and China has declared robotics as a core industry.³⁹ As the United States reshores certain strategic supply chains and increases its reliance on autonomous systems, continued access to cutting-edge robotics will be a national security imperative.

Action for the National Institute of Standards and Technology:

- **Incentivize the development of world-class software platforms for robotic systems by U.S. firms.**
 - o By designing the software platforms upon which core robotic capabilities are built, U.S. firms will be well-positioned to shape the next wave of industrialization. The U.S. government should expand collaboration with industry on basic R&D, set international standards, and share data pertaining to robotic system development by expanding upon the work of the Intelligent Systems Division at the National Institute of Standards and Technology (NIST).⁴⁰ The U.S. government should also incentivize the early adoption of robotic systems across the public and private sectors by creating markets in areas ripe for automation.⁴¹ These efforts will yield valuable data and experience in scaling automation and facilitate the application of robotics to adjacent sectors. A multipronged approach along these lines will position U.S. industry to compete more effectively in the market for robotic systems software, a strategically important area that is compatible with existing U.S. strengths.

Advanced Manufacturing

Recommendation: Accelerate Additive Manufacturing Production of Legacy Parts Across the Department of Defense

Recommendation

The ability to manufacture high-tech products domestically is critical to a nation's security and its economic productivity. The United States must strive to develop manufacturing capabilities in industries that are essential to crisis response or that would take too long to bring online in the event of a protracted conflict.⁴² Innovation also benefits from the co-location of firms engaged in technological design and those that produce finished products, which enables rapid feedback and continuous iteration on product design.⁴³ This link is particularly important in the defense sector, where communication between

researchers, designers, and manufacturers can help quickly transition a technology from the lab to the field. However, the United States has relinquished manufacturing leadership in high-tech industries that employ highly skilled workers to high-wage nations like Germany and Japan.⁴⁴ Meanwhile, China and other lower-wage nations are moving up the value chain from low-value manufacturing processes, such as assembly, to more sophisticated techniques.⁴⁵ Although the supply chain disruptions resulting from the COVID-19 pandemic may prompt the return of some manufacturing to the United States, the broader trend of offshoring the manufacturing of next-generation technologies appears likely to continue unless the U.S. government takes appropriate action.⁴⁶

Action for the Department of Defense:

- **Accelerate additive manufacturing of legacy parts across the Department of Defense.**
 - Additive manufacturing and 3D printing have the potential to transform the manufacturing industry by enabling the rapid production of complex objects on demand and at the point of need.⁴⁷ Although existing 3D printers cannot match the quality of advanced traditional techniques, AI has shown the potential to significantly improve the accuracy of 3D printing.⁴⁸ The DoD should proactively support the improvement of 3D printing by identifying all legacy parts in active weapon systems suited to production by additive manufacturing and 3D printers and commit to doing so by 2025.⁴⁹

Energy Systems

Recommendation

Recommendation: Develop and Domestically Manufacture Energy Storage Technologies to Meet U.S. Market Demand by 2030

Cheap and reliable access to energy is critical to U.S. national security. Although the United States is at the forefront of the exploration, extraction, and processing of oil and gas and possesses significant domestic reserves, China is by far and away the leading producer of renewable energy and is investing heavily in advanced energy storage technologies, such as batteries and their constituent materials.⁵⁰ As the cost of intermittent renewable sources continues to fall, the United States must commit to developing and deploying the next generation of energy storage devices, from long-duration stationary applications to battery packs for electric vehicles.

Action for Congress:

- **Fund the Department of Energy's initiative to develop and domestically manufacture energy storage technologies to meet U.S. market demand by 2030.**
 - Improving the cost and energy density of storage technologies will drive progress in sectors ranging from electric vehicles to distributed energy generation. The Department of Energy (DoE) has set the ambitious goal of developing and domestically manufacturing storage technologies capable of meeting the entirety

of U.S. market demand by 2030.⁵¹ Congress should fully fund the federal R&D needed to achieve the DoE's Energy Storage Grand Challenge roadmap by 2030 and establish appropriate incentives for the commercialization of the resulting technologies.⁵²

Further Consideration of Additional Technologies and Conclusion

While the Commission believes the eight emerging technologies discussed above and elsewhere in this report—AI, microelectronics, biotechnology, quantum computing, 5G telecommunications, autonomy and robotics, advanced manufacturing, and energy systems—will be crucial to future national competitiveness, this list is by no means exhaustive. Other emerging technologies and platforms—everything from digital currencies and other types of financial technology to space systems—will likely also play a major role in the U.S. economy and its national security moving forward. And there are undoubtedly technologies that have yet to be created which, in the near future, will have transformative effects on the lives and security of American citizens.

We are at the beginning of a new era, in which technologies not only are the principal driver of global markets and geopolitics, but they also advance and emerge faster than ever before. As the speed of technological development accelerates and an increasing number of technologies have dual-use applications, techno-national security threats will continue to multiply. To meet this challenge, the U.S. government must continually assess new technological advancements to determine their potential to disrupt industries, change economies, and transform national security.

The process of technology horizon-scanning, forecasting, and proactively crafting policies to address upcoming national security threats related to emerging technologies must become an ingrained component of the U.S. national security process. Doing so is not only essential, but also urgent. If the U.S. government waits to adapt to this new reality until a subsequent commission makes a similar recommendation, it will likely be playing technological catch-up from a position of national security weakness. As existing technologies evolve and new ones emerge, the relationship between technology and national security will only grow stronger, and the need for the United States to maintain overall technical leadership will only increase.

Blueprint for Action: Chapter 16 - Endnotes

¹ The Commission identified these as essential to overall U.S. technological leadership in its 2019 Interim Report. See *Interim Report*, NSCAI at 31 (Nov. 2019), <https://www.nscai.gov/previous-reports/>.

² *Summary of the 2018 National Defense Strategy of the United States of America*, U.S. Department of Defense at 3 (2018), <https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf>.

³ *Fiscal Year 2019 Industrial Capabilities: Report to Congress*, U.S. Department of Defense at 132 (June 23, 2020), <https://www.businessdefense.gov/Portals/51/Documents/Resources/USA000954-20%20RPT%20Subj%20FY19%20ICR%2007092020.pdf?ver=2020-07-10-124452-180>.

⁴ 83 Fed. Reg. 58201, *Review of Controls for Certain Technologies*, U.S. Department of Commerce: Bureau of Industry and Security (Nov. 19, 2018), <https://www.federalregister.gov/documents/2018/11/19/2018-25221/review-of-controls-for-certain-emerging-technologies>.

⁵ *Recommendations for Strengthening American Leadership in Industries of the Future*, President's Council of Advisors on Science and Technology (June 2020), https://science.osti.gov/-/media/_/pdf/about/pcast/202006/PCAST_June_2020_Report.pdf.

⁶ S. 3832, 116th Cong. (2020), <https://www.congress.gov/bill/116th-congress/senate-bill/3832>.

⁷ *National Strategy for Critical and Emerging Technologies*, The White House at A-1 (Oct. 2020), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2020/10/National-Strategy-for-CET.pdf>.

⁸ Donald J. Trump, *Executive Order 13922: Delegating Authority Under the Defense Production Act to the Chief Executive Officer of the United States International Development Finance Corporation To Respond to the COVID-19 Outbreak*, The White House (May 14, 2020), <https://www.federalregister.gov/documents/2020/05/19/2020-10953/delegating-authority-under-the-defense-production-act-to-the-chief-executive-officer-of-the-united>.

⁹ For comparison, the Chinese government provided approximately \$117 million in initial funding to the China National GeneBank for its construction and creation. See Zhuang Pinghui, *China Opens First National Gene Bank, Aiming to House Hundreds of Millions of Samples*, South China Morning Post (Sept. 22, 2016) <https://www.scmp.com/news/china/article/2021623/chinas-noahs-ark-first-national-gene-bank-opens-shenzhen>.

¹⁰ Elsa Kania, *Minds at War: China's Pursuit of Military Advantage Through Cognitive Science and Biotechnology*, Prism (Jan. 2020), https://ndupress.ndu.edu/Portals/68/Documents/prism/prism_8-3/prism_8-3_Kania_82-101.pdf.

¹¹ This recommendation is included in Chapter 1 of this report.

¹² *National Biodefense Strategy*, The White House (2018), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2018/09/National-Biodefense-Strategy.pdf>.

¹³ See Chapter 9 of this report for additional details on the proposed Technology Competitiveness Council.

¹⁴ See, e.g., *BGI & US Collaborate on Precision Medicine Development*, UW Medicine (May 10, 2016), <https://newsroom.uw.edu/story/bgi-uw-collaborate-precision-medicine-development>.

¹⁵ In 2010, BGI received a \$1.5 billion loan from the state-run China Development Bank. The precise extent of government subsidies to BGI are unknown, but likely substantial. See Kirsty Needham, *Special Report: COVID Opens New Doors for China's Gene Giant*, Reuters (Aug. 5, 2020), <https://www.reuters.com/article/us-health-%20coronavirus-bgi-specialreport/special-report-covid-opens-new-doors-for-chinas-gene-giant-idUSKCN2511CE>; see also Antonio Regalado, *China's BGI Says It Can Sequence a Genome for Just \$100*, MIT Technology Review (Feb. 26, 2020), <https://www.technologyreview.com/2020/02/26/905658/china-bgi-100-dollar-genome/>.

¹⁶ *China National Genebank Officially Opens*, BGI (Sept. 22, 2016), <https://www.bgi.com/us/company/careers/china-national-genebank-officially-opens/>.

¹⁷ See Kirsty Needham, *Special Report: COVID Opens New Doors for China's Gene Giant*, Reuters (Aug. 5, 2020), <https://www.reuters.com/article/us-health-coronavirus-bgi-specialreport/special-report-covid-opens-new-doors-for-chinas-gene-giant-idUSKCN2511CE>; see also Jeanne Whalen & Elizabeth Dwoskin, *California Rejected Chinese Company's Push to Help with Coronavirus Testing. Was That the Right Move?* Washington Post (July 2, 2020), <https://www.washingtonpost.com/business/2020/07/02/china-bgi-california-testing/>.

¹⁸ See Chapters 9 and 15 of this report for additional information on cooperation on issues at the intersection of AI and global health.

¹⁹ Pontus Vikstål, et al., *Applying the Quantum Approximate Optimization Algorithm to the Tail-Assignment Problem*, Physical Review Applied Vol. 14, Iss. 3 (Sept. 3, 2020), <https://doi.org/10.1103/PhysRevApplied.14.034009>; He Ma, et al., *Quantum Simulations of Materials on Near-Term Quantum Computers*, npj Computational Materials (July 2, 2020), <https://doi.org/10.1038/s41524-020-00353-z>.

²⁰ C.L. Degen, et al., *Quantum Sensing*, arXiv (June 7, 2017), <https://arxiv.org/pdf/1611.02427.pdf>; Juan Yin, et al., *Entanglement-Based Secure Quantum Cryptography over 1,120 Kilometres*, Nature 582, 501-505 (June 15, 2020), <https://doi.org/10.1038/s41586-020-2401-y>.

²¹ In December 2020, a team of researchers in China demonstrated quantum advantage on a photonic quantum computer. See Han-Sen Zhong, et al., *Quantum Computational Advantage Using Photons*, Science (Dec. 18, 2020), <https://science.sciencemag.org/content/370/6523/1460>.

²² Elsa B. Kania & John Costello, *Quantum Hegemony? China's Ambitions and the Challenge to U.S. Innovation Leadership*, CNAS (Sept. 12, 2018), <https://www.cnas.org/publications/reports/quantum-hegemony>.

²³ For additional information and NSCAI views on quantum computing, AI, and national security, see *Interim Report* and *Third Quarter Recommendations*, NSCAI at 154-163 (Oct. 2020), <https://www.nscai.gov/previous-reports/>.

²⁴ *Applications of Quantum Technologies: Executive Summary*, Defense Science Board at C-1 (Oct. 2019), https://dsb.cto.mil/reports/2010s/DSB_QuantumTechnologies_Executive%20Summary_10.23.2019_SR.pdf.

²⁵ Dario Gil, *How to Ensure the U.S.'s Quantum Future*, Scientific American (Aug. 20, 2020), <https://www.scientificamerican.com/article/how-to-ensure-the-uss-quantum-future/>.

²⁶ The bipartisan National Quantum Initiative Act of 2018 prompted a series of steps that establish quantum computing as a strategic priority for the United States. The Department of Energy (DoE) announced \$625 million to establish five Quantum Information Science research centers over five years led by the national laboratories. The National Science Foundation (NSF) announced \$75 million to create three Quantum Leap Challenge Institutes over the same period. Lastly, the President's FY 2021 Budget recommended doubling federal investment levels in quantum technologies by 2022. This continuing investment is necessary to determine the full potential of quantum computing and maintain the United States' position of leadership in next-generation computer hardware. For more details, see Pub. L. 115-368, National Quantum Initiative Act, 115th Cong. (2018), <https://www.congress.gov/bill/115th-congress/house-bill/6227>; *Department of Energy Announces \$625 Million for New Quantum Centers*, U.S. Department of Energy (Jan. 10, 2020), <https://www.energy.gov/articles/department-energy-announces-625-million-new-quantum-centers>; *NSF Establishes 3 New Institutes to Address Critical Challenges in Quantum Information Science*, National Science Foundation (July 21, 2020), https://www.nsf.gov/news/special_reports/announcements/072120.jsp; *Recommendations for Strengthening American Leadership in Industries of the Future*, The President's Council of Advisors on Science and Technology at 13 (June 2020), https://science.osti.gov/-/media/_pdf/about/pcast/202006/PCAST_June_2020_Report.pdf?la=en&hash=019A4F17C79FDEE5005C51D3D6CAC81FB31E3ABC.

²⁷ John Preskill, *Quantum Computing in the NISQ Era and Beyond* arXiv at 4, 14 (July 30, 2018), <https://arxiv.org/pdf/1801.00862.pdf>.

Blueprint for Action: Chapter 16 - Endnotes

²⁸ The Department of Energy and the Air Force offer access to commercial quantum capabilities, but this access is not widespread, nor is it focused on hybrid quantum-classical software development. See Adrian Cho, *After Years of Avoidance, Department of Energy Joins Quest to Develop Quantum Computers*, *Science* (Jan. 10, 2018), <https://www.sciencemag.org/news/2018/01/after-years-avoidance-department-energy-joins-quest-develop-quantum-computers>; *Air Force Research Laboratory to Join IBM Q Network as First DOD-led IBM Q Hub*, Wright-Patterson AFB (Aug. 2, 2019), <https://www.wpafb.af.mil/News/Article-Display/Article/1924271/air-force-research-laboratory-to-join-ibm-q-network-as-first-dod-led-ibm-q-hub/>.

²⁹ See *First Quarter Recommendations*, NSCAI at 12-13 (March 2020), <https://www.nscai.gov/previous-reports/>. In the FY 2021 NDAA, Congress took the first step toward implementing the Commission's First Quarter recommendations by creating a task force to develop a road map for the NAIRR. See Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021). In Chapter 11 of this report, the Commission urges Congress to authorize and appropriate the funds necessary to carry out the task force's road map immediately.

³⁰ Dan Strumpf, *U.S. vs. China in 5G: The Battle Isn't Even Close*, *Wall Street Journal* (Nov. 9, 2020), <https://www.wsj.com/articles/u-s-vs-china-in-5g-the-battle-isnt-even-close-11604959200>.

³¹ Dana Deasy, *Department of Defense Statement on Mid-Band Spectrum*, U.S. Department of Defense (Aug. 10, 2020), <https://www.defense.gov/Newsroom/Speeches/Speech/Article/2307288/department-of-defense-statement-on-mid-band-spectrum>.

³² Milo Medin & Gilman Louie, *The 5G Ecosystem: Risks & Opportunities for DoD*, Defense Innovation Board (April 3, 2019), https://media.defense.gov/2019/Apr/03/2002109302/-1/-1/0/DIB_5G_STUDY_04.03.19.PDF.

³³ *3.5 GHz Band Overview*, U.S. Federal Communications Commission (April 23, 2020), <https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-overview>.

³⁴ Edward Drocella, et al., *Technical Feasibility of Sharing Federal Spectrum with Future Commercial Operations in the 3450-3550 MHz Band*, NTIA (Jan. 27, 2020), <https://www.ntia.gov/report/2020/technical-feasibility-sharing-federal-spectrum-future-commercial-operations-3450-3550>.

³⁵ See *Public Notice: Auction of Priority Access Licenses in the 3550-3650 MHz Band Closes*, FCC (Sept. 2, 2020), <https://docs.fcc.gov/public/attachments/DA-20-1009A1.pdf>.

³⁶ In April 2019, the Defense Innovation Board issued a report which argued that the status quo of spectrum allocation is unsustainable and DoD must expand its sub-6GHz spectrum-sharing operations to enable the United States to compete with China in 5G. Milo Medin & Gilman Louie, *The 5G Ecosystem: Risks & Opportunities for DoD*, Defense Innovation Board (April 3, 2019), https://media.defense.gov/2019/Apr/03/2002109302/-1/-1/0/DIB_5G_STUDY_04.03.19.PDF. For additional information and views of the Commission on 5G, see *First Quarter Recommendations*, NSCAI at 54-57 (March 2020), <https://www.nscai.gov/previous-reports/>.

³⁷ Recommendations to ensure U.S. leadership in biotechnology, quantum computing, and 5G telecommunications can be found above. Recommendations pertaining to semiconductors and advanced computer hardware are covered in Chapter 13 of this report.

³⁸ *Summer Study on Autonomy*, Defense Science Board (June 2016), <https://dsb.cto.mil/reports/2010s/DSBSS15.pdf>.

³⁹ China owns one-fifth of the global supply of industrial robots and sought to have 45% of its high-end robots be produced domestically by the end of 2020. See Johnny Williamson, *How Nations Around the World Are Investing in Robotics Research*, *The Manufacturer* (June 10, 2020), <https://www.themanufacturer.com/articles/how-nations-around-the-world-are-investing-in-robotics-research/>.

⁴⁰ *Intelligent Systems Division*, NIST (last accessed Feb. 2, 2021), <https://www.nist.gov/el/intelligent-systems-division-73500>.

⁴¹ For example, the U.S. Postal Service could scale its Autonomous Mobile Robot pilot program from 25 sorting facilities to all sorting facilities by 2025. *Autonomous Mobile Robots and the Postal Service*, U.S. Postal Service Office of Inspector General (April 9, 2018), <https://www.uspsoidg.gov/sites/default/files/document-library-files/2019/RARC-WP-18-006.pdf>.

⁴² An Executive Order 13806 report identifies 10 manufacturing risk archetypes. See *Assessing and Strengthening the Manufacturing and Defense Industrial Base and Supply Chain Resiliency of the United States*, U.S. Department of Defense (Sept. 2018), <https://media.defense.gov/2018/oct/05/2002048904/-1/-1/1/assessing-and-strengthening-the-manufacturing-and%20defense-industrial-base-and-supply-chain-resiliency.pdf>.

⁴³ *Strategy for American Leadership in Advanced Manufacturing*, National Science and Technology Council (Oct. 2018), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2018/10/Advanced-Manufacturing-Strategic-Plan-2018.pdf>; Gregory Tasse, *Rationales and Mechanisms for Revitalizing US Manufacturing R&D Strategies*, NIST (Jan. 29, 2010), https://www.nist.gov/system/files/documents/2017/05/09/manufacturing_strategy_paper_0.pdf.

⁴⁴ *Report to the President on Ensuring American Leadership in Advanced Manufacturing*, President's Council of Advisors on Science and Technology (June 2011), <https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/pcast-advanced-manufacturing-june2011.pdf>; *Advanced Manufacturing: A Snapshot of Priority Technology Areas Across the Federal Government*, National Science and Technology Council (April 2016), https://www.manufacturing.gov/sites/default/files/2018-01/nstc_sam_technology_areas_snapshot.pdf.

⁴⁵ *Report to the President on Ensuring American Leadership in Advanced Manufacturing*, President's Council of Advisors on Science and Technology at 3 (June 2011), <https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/pcast-advanced-manufacturing-june2011.pdf>.

⁴⁶ *Advanced Manufacturing: Innovation Institutes Have Demonstrated Initial Accomplishments, but Challenges Remain in Measuring Performance and Ensuring Sustainability*, GAO-19-409 at 1 (May 23, 2019), <https://www.gao.gov/assets/700/699310.pdf>.

⁴⁷ *Audit of the DoD's Use of Additive Manufacturing for Sustainment Parts*, U.S. Department of Defense Inspector General (Oct. 17, 2019), <https://media.defense.gov/2019/Oct/21/2002197659/-1/-1/1/DODIG-2020-003.PDF>.

⁴⁸ Mark Anderson, *3D Print Jobs Are More Accurate With Machine Learning*, IEEE Spectrum (Feb. 19, 2020), <https://spectrum.ieee.org/tech-talk/artificial-intelligence/machine-learning/3d-print-jobs-news-accurate-machine-learning>.

⁴⁹ For instance, in August 2020, the Department of Defense printed the first metal part for a B-52 jet engine. See Kyle Mizokami, *The Old-School Engine That Powers the B-52 Gets a 3D-Printed Upgrade*, Popular Mechanics (Aug. 10, 2020), <https://www.popularmechanics.com/military/aviation/a33535790/air-force-3d-print-metal-part-turbofan-engine/>.

⁵⁰ Robert Rapier, *Ten Countries That Dominate Fossil Fuel Production*, Forbes (July 14, 2019), <https://www.forbes.com/sites/rrapier/2019/07/14/ten-countries-that-dominate-fossil-fuel-production>; *Country Rankings*, International Renewable Energy Agency, (last accessed Feb. 2, 2021), <https://www.irena.org/Statistics/View-Data-by-Topic/Capacity-and-Generation/Country-Rankings>.

⁵¹ *Energy Storage*, U.S. Department of Energy (last accessed Feb. 2, 2021), <https://www.energy.gov/oe/energy-storage>.

⁵² *Energy Storage Grand Challenge: Roadmap*, U.S. Department of Energy (Dec. 2020), <https://www.energy.gov/sites/prod/files/2020/12/f81/Energy%20Storage%20Grand%20Challenge%20Roadmap.pdf>.

Appendices

Appendix A: Technical Glossary	601
Appendix B: Acronyms Found in This Report	617
Appendix C: Key Considerations for the Responsible Development and Fielding of Artificial Intelligence (Abridged)	633
Appendix D: Draft Legislative Language	663
Appendix E: Funding Recommendation Table	729
Appendix F: Commissioner Bios	741
Appendix G: Commission Staff and Contributors	749

Appendix A: Technical Glossary

3D Chip Stacking: The process of building integrated circuits with both horizontal and vertical interconnections between transistors. This brings elements of the chip physically closer together, increasing density and allowing for greater performance (i.e., speed) at lower power levels and at a smaller footprint than comparable two-dimensional devices, which only feature horizontal interconnects.

Additive Manufacturing: A computer-controlled process in which successive layers of material are deposited to create a part that matches a 3D design.

Adversarial Machine Learning: A broad collection of techniques used to exploit vulnerabilities across the entire machine learning stack and lifecycle. Adversaries may target the data sets, algorithms, or models that an ML system uses in order to deceive and manipulate their calculations, steal data appearing in training sets, compromise their operation, and render them ineffective.¹ Adversarial AI may be used as a phrase that broadens the considerations to attacks on AI systems, including approaches that are less dependent on data and machine learning.

Agile: A philosophy and methodology used to describe the continuous, iterative process to develop and deliver software and other digital technologies. User requirements and feedback inform incremental development and delivery by developers.²

AI Assurance: The defensive science of protecting AI applications from attack or malfunction.

AI Digital Ecosystem: A technology stack driving the development, testing, fielding, and continuous update of AI-powered applications. The ecosystem is managed as a multi-layer collection of shared AI essential building blocks (e.g., data, algorithms, tools, and trained AI models) accessed through common interfaces.

AI Governance: The actions to ensure stakeholder needs, conditions, and options are evaluated to determine balanced, agreed-upon enterprise objectives; setting direction through prioritization and decision-making; and monitoring performance and compliance against agreed-upon directions and objectives.³ AI governance may include policies on the nature of AI applications developed and deployed versus those limited or withheld.

AI Lifecycle: The steps for managing the lifespan of an AI system: 1) Specify the system's objective. 2) Build a model. 3) Test the AI system. 4) Deploy and maintain the AI system. 5) Engage in a feedback loop with continuous training and updates.⁴

AI Stack: AI can be envisioned as a stack of interrelated elements: talent, data, hardware, algorithms, applications, and integration.⁵

Algorithm: A series of step-by-step instructions or calculations to solve an instance of a problem. There are fundamentally two ways that algorithms are implemented by AI: explicit engineering of the algorithm (e.g., in symbolic reasoning and expert systems) or by machine learning, where the algorithm is derived from data or feedback from interactions.

Anonymization: Also referred to as data de-identification, this is the process of removing or replacing with synthetic values any identifiable information in data. This is intended to make it impossible to derive insights on any specific individual in the data while remaining useful for the intended use of the data.⁶ (See de-anonymization.)

Application Programming Interfaces (APIs): Programming tools for describing how one program can access the functionality of another⁷ while hiding the implementation details inside each program.

Application-Specific Integrated Circuit (ASIC): A chipset custom designed to perform a particular task. ASICs could provide significant performance gains over generic chips but are inflexible in their functions compared to central processing units.

Architecture: A set of values, constraints, guidance, and practices that support the active evolution of the planning, designing, and construction of a system. The approach evolves over time, while simultaneously supporting the needs of current customers.⁸ Architecture can refer to sets of components in a computing system and their operational interrelationships as well as other important configurations such as the architecture of a neural network, which captures the patterns of connectivity within and between layers of units in the network model.

Artificial General Intelligence (AGI): A phrase that has been used to capture the possibility of developing more general AI capabilities, in distinction to the typically narrow capabilities of AI systems that have been developed to date. Some use the term to refer to the prospect of achieving more human-like intelligence, developing AI systems with the ability to perform many of the intellectual tasks that humans are capable of doing, or developing systems that might employ a wide range of skills across multiple domains of expertise.

Artificial Intelligence (AI): The ability of a computer system to solve problems and to perform tasks that have traditionally required human intelligence to solve.

Auditability: A characteristic of an AI system in which its software and documentation can be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations.

Augmented Reality: Enhanced digital content, spanning visual, auditory, or tactile information, overlaid onto the physical world.⁹

Authorization to Operate (ATO): The official management decision given by a senior organizational official to authorize operation of an information system and to explicitly accept the risk to organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, and other organizations based on the implementation of an agreed-upon set of security controls.¹⁰

Automation Bias: An unjustified degree of reliance on automated systems or their outcomes.

Autonomous: A system with functions capable of operating without direct human control.

Biological Sensors (Biosensors): Devices used to detect the presence or concentration of a biological analyte, such as a biomolecule, a biological structure, or a microorganism. Biosensors consist of three parts: a component that recognizes the analyte and produces a signal, a signal transducer, and a reader device.¹¹

Biometric Technologies: Technologies that leverage physical or behavioral human characteristics that can be used to digitally identify a person and grant access to systems, devices, or data, such as face, voice, and gait recognition.¹²

Black Box: The nature of some AI techniques whereby the inferential operations are complex, hidden, or otherwise opaque to their developers and end users in terms of providing an understanding of how classifications, recommendations, or actions are generated and what overall performance will be.

Carbon Nanotubes: Nano-scale structures that can be used to make transistors and could potentially replace silicon transistors in the future. Compared to existing silicon transistors, carbon nanotube transistors are both capable of being shrunk to a smaller size and more amenable to being stacked in three dimensions (see 3D chip stacking).

Cloud Computing: The act of running software within information technology environments that abstract, pool, and share scalable resources across a network.¹³

Cloud Infrastructure: The components needed for cloud computing, which include hardware, abstracted resources, storage, and network resources.¹⁴

Commonsense Reasoning: The process of forming a conclusion based on the basic ability to perceive, understand, and judge things that are shared by (“common to”) most people and can reasonably be expected without need for debate.¹⁵ Endowing computing systems with the commonsense knowledge of humans has been found to be a difficult and standing AI challenge.

Computational Thinking: The thought processes involved in formulating problems so their solutions can be represented as computational steps and algorithms.¹⁶

Computer Vision: The digital process of perceiving and learning visual tasks in order to interpret and understand the world through cameras and sensors.¹⁷

Continuous Delivery: A process that builds on continuous integration by taking the step of orchestrating multiple builds, coordinating different levels of automated testing, and moving the code into a production environment in a process that is as automated as possible.¹⁸

Continuous Integration: A process that aims to minimize the duration and effort required by “each” integration episode and deliver at any moment a product version suitable for release. In practice, this requires an integration procedure that is reproducible and mostly automated. This is achieved through version control tools, team policies, and conventions.¹⁹

Data Architecture: The structure of an organization’s logical and physical data assets and data management resources.²⁰

Data Privacy: The right of an individual or group to maintain control over, and the confidentiality of, information about themselves.²¹

Data Protection: The practice of safeguarding information from unauthorized access, use, disclosure, disruption, modification, or destruction, to provide confidentiality, integrity, and availability.²²

De-anonymization: Matching anonymous data (also known as de-identified data) with publicly available information, or auxiliary data, in order to discover the individual to whom the data belong.²³ (See anonymization.)

Deepfake: Computer-generated video or audio (particularly of humans) so sophisticated that it is difficult to distinguish from reality.²⁴ Deepfakes have also been referred to as synthetic media.

Deep Learning: A machine learning implementation technique that exploits large quantities of data, or feedback from interactions with a simulation or the environment, as training sets for a network with multiple hidden layers, called a deep neural network, often employing

an iterative optimization technique called gradient descent, to tune large numbers of parameters that describe weights given to connections among units.²⁵

Deep Neural Networks (DNN): A deep learning architecture that is trained on data or feedback, generating outputs, calculating errors, and adjusting its internal parameters. The process is repeated possibly hundreds of thousands of times until the network achieves an acceptable level of performance. It has proved to be an effective technique for image classification, object detection, speech recognition, some kinds of game-playing, and natural language processing—problems that challenged researchers for decades. By learning from data, DNNs can solve some problems much more effectively and also solve problems that were never solvable before.²⁶

Deployed AI: AI that has been fielded for its intended purpose within its relevant operational environment.

DevSecOps: Enhanced engineering practices that improve the lead time and frequency of delivery outcomes, promoting a more cohesive collaboration between development, security, and operations teams as they work toward continuous integration and delivery.²⁷

Differential Privacy: A criterion for a strong, mathematical definition of privacy in the context of statistical and machine learning analysis used to enable the collection, analysis, and sharing of a broad range of statistical estimates, such as averages, contingency tables, and synthetic data, based on personal data while protecting the privacy of the individuals in the data.²⁸

Digital Ecosystem: The stakeholders, systems, tools, and enabling environments that together empower people and communities to use digital technology to gain access to services, engage with each other, and pursue missional opportunities.²⁹

Digital Infrastructure: The foundational components that enable digital technologies and services. Examples of digital infrastructure include fiber-optic cables, cell towers, satellites, data centers, software platforms, and end-user devices.³⁰

Distributed System: A system whose components are located on different networked computers, which communicate and coordinate their actions by passing messages to one another in order to appear as a single system to the end user.³¹

Domain-Specific Hardware Architectures: Hardware that is specifically designed to fulfill certain narrow functions, seeking performance gains through specialization.

Edge Computing: A distributed-computing paradigm that brings computation and data storage closer to the location where it is needed (i.e., the network edge where smart sensors, devices, and systems reside along with points of human interaction) to improve response times and save bandwidth.³²

Expert System: A computer system emulating the decision-making ability of a human expert through the use of reasoning, leveraging an encoding of domain-specific knowledge most commonly represented by sets of if-then rules rather than procedural code.³³ The term “expert system” was used largely during the 1970s and '80s amidst great enthusiasm about the power and promise of rule-based systems that relied on a “knowledge base” of domain-specific rules and rule-chaining procedures that map observations to conclusions or recommendations.

Explainability: A characteristic of an AI system in which there is provision of accompanying evidence or reasons for system output in a manner that is meaningful or understandable to individual users (as well as to developers and auditors) and reflects the system's process for generating the output (e.g., what alternatives were considered, but not proposed, and why not).³⁴

False Negative: An example in which the predictive model mistakenly classifies an item as in the negative class. For example, a false negative describes the situation in which a junk-email model specifies that a particular email message is not spam (the negative class) when the email message actually is spam, leading to the frustration of the junk message appearing in an end user's inbox.³⁵ In a higher-stakes example, a false negative captures the case in which a medical diagnostic model misses identifying a disease that is present in a patient.

False Positive: An example in which the predictive model mistakenly classifies an item as in the positive class. For example, the model inferred that a particular email message was spam (the positive class), but that email message was actually not spam, leading to delays in an end user reading a potentially important message.³⁶ In a higher-stakes situation, a false positive describes the situation in which a disease is diagnosed as present when the disease is not present, potentially leading to unnecessary and costly treatments.

Federated Data Repository: A virtual data repository that links data from distributed sources (e.g., other repositories), providing a common access portal for finding and accessing data.

Field-Programmable Gate Array (FPGA): An integrated circuit featuring reconfigurable interconnects that can be programmed by the user to be customized for specific functions after it is manufactured. FPGAs feature greater flexibility than ASICs, but at a cost to performance.

Gallium Nitride: An alternative material to silicon for transistors. Gallium nitride transistors feature higher electron mobility than silicon and are capable of faster switching speed, higher thermal conductivity, and lower on-resistance than comparable silicon solutions.

Generative Adversarial Networks (GANs): An approach to training AI models useful for applications like data synthesis, augmentation, and compression where two neural networks are trained in tandem: one is designed to be a generative network (the forger) and the other a discriminative network (the forgery detector). The objective is for each network to train and better itself off the other, reducing the need for big labeled training data.³⁷

Graphics Processing Unit (GPU): A specialized chip capable of highly parallel processing. GPUs are well-suited for running machine learning and deep learning algorithms. GPUs were first developed for efficient parallel processing of arrays of values used in computer graphics. Modern-day GPUs are designed to be optimized for machine learning.

High-Performance Computing (HPC): Developing, deploying, and operating very high-capacity computers (along with the requisite software, hardware, facilities, and underpinning infrastructure) to advance the computational upper limits of resolution, dimensionality, and complexity.³⁸

Homomorphic Encryption: A technique that allows computation to be performed directly on encrypted data without requiring access to a secret key. The result of such a computation remains in encrypted form and can at a later point be revealed by the owner of the secret key.³⁹

Human-Machine Teaming (or Human-AI Teaming): The ability of humans and AI systems to work together to undertake complex, evolving tasks in a variety of environments with seamless handoff both ways between human and AI team members. Areas of effort include developing effective policies for controlling human and machine initiatives,⁴⁰ computing methods that ideally complement people,⁴¹ methods that optimize goals of teamwork, and designs⁴² that enhance human-AI interaction.

Information Operations: The tactics, techniques, and procedures employed in both the offensive and defensive use of information to pursue a competitive advantage.⁴³

Internet of Things (IoT): A global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies.⁴⁴

Intelligent Sensing: Utilizing advanced signal processing techniques, data fusion techniques, intelligent algorithms, and AI concepts to better understand sensor data for better integration of sensors and better feature extraction, leading to actionable knowledge that can be used in smart sensing applications.⁴⁵

Interpretability: The ability to understand the value and accuracy of system output. Interpretability refers to the extent to which a cause and effect can be observed within

a system or to which what is going to happen given a change in input or algorithmic parameters can be predicted. Interpretability complements explainability.⁴⁶

Legacy Systems: Outdated systems still in operation that are hard to maintain owing to shortage of skill sets and obsolete architecture.⁴⁷

Machine Learning (ML): The study or the application of computer algorithms that improve automatically through experience.⁴⁸ Machine learning algorithms build a model based on training data in order to perform a specific task, like aiding in prediction or decision-making processes, without necessarily being explicitly programmed to do so.

Microelectronics: A subfield of electronics involving small components such as transistors, capacitors, and resistors. These components are packaged together to form the integrated circuits that are used to perform computations.

MLOps: Enhanced engineering practices that combine ML model development and ML model operations technologies to support continuous integration and delivery of ML-based solutions.⁴⁹

Modeling and Simulation: Modeling the physical world to support the study, optimization, and testing of operations through simulation without interfering or interrupting ongoing processes. Modeling and simulation can be used to train AI systems, and AI technologies can be used to enhance modeling and simulation.

Multi-Party Federated Learning: An ML setting where many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., service provider) while keeping the training data decentralized. It can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized ML and data science approaches.⁵⁰ However, it does introduce new attack vectors that must be addressed.⁵¹

Multi-Source Data: Data obtained and aggregated from different origins.

Multimodal Data: Data comprising several signal or communication types, such as speech and body gestures during human-to-human communication.

Natural Language Processing: The ability of a machine to process, analyze, and mimic human language, either spoken or written.

Natural Language Understanding: The ability of a machine to represent and act on the meaning that a language expresses utilizing language semantically rather than statistically.

Neuromorphic Computing: Computing that mimics the human brain or neural network.⁵²

Object Recognition: The algorithmic process of finding objects in the real world from an image, typically using object models which are known a priori.⁵³

One Shot (or Few Shot) Learning: An approach to machine learning that leverages existing knowledge to enable learning in some applications (e.g., object recognition) on a few non-repeated examples, with the system rapidly learning similarities and dissimilarities between the training examples.⁵⁴

Open Knowledge Network (OKN): A vision to create an open knowledge graph of all known entities and their relationships, ranging from the macro (e.g., have there been unusual clusters of earthquakes in the U.S. in the past six months?) to the micro (e.g., what is the best combination of chemotherapeutic drugs for a 56-year-old female with stage 3 brain cancer?). OKN is meant to be an inclusive, open, community activity resulting in a knowledge infrastructure that could facilitate and empower a host of applications and open new research avenues, including how to create trustworthy knowledge networks/graphs.⁵⁵

Packaging: The final stage of the semiconductor fabrication process, in which a chip is placed in its protective case. For many years packaging was a low-value element of the semiconductor design process. However, advanced packaging techniques are enabling sophisticated new chip designs using processes such as 3D stacking, heterogeneous integration, and modular chiplets to create more complex and sophisticated semiconductors.

Pattern Recognition: The field concerned with the automatic discovery of regularities in data through the use of computer algorithms, with the use of these regularities to take actions such as classifying the data into different categories.⁵⁶

Planning and Optimization: Determining necessary steps to complete a series of tasks, which can save time and money and improve safety.

Platform Environment: Provides an application developer or user secured access to resources and tools (e.g., workflows, data, software tools, storage, and compute) on which applications can be developed or run.

Polymorphic Malware: A type of malware that constantly changes its identifiable features (i.e., signatures) in order to evade detection. Many of the common forms of malware can be polymorphic, including viruses, worms, bots, trojans, or keyloggers.⁵⁷

Precision: A metric for classification models. Precision identifies the frequency with which a model was correct when classifying the positive class. It answers the question “How many selected positive items are true positive?”—for example, the percentage of messages flagged as spam that actually are spam.⁵⁸

Prediction: Forecasting quantitative or qualitative outputs through function approximation, applied on input data or measurements.⁵⁹

Prior Art: The worldwide scientific and technical knowledge by which an invention is evaluated to determine if it is new.

Pseudonymization: A data management technique to strip identifiers linking data to an individual. Concern exists that such data could still be linked with other data that allows for a person's identity to be rediscovered.

PyTorch: A free and open-source software library for training neural networks and other machine learning architectures, initially developed by Facebook AI Research.

Quantum Computer: A machine that relies on the properties of quantum mechanics to perform computations. Quantum computers encode information in *qubits*, which can exist in a linear combination of two states. These states can be physically realized in a number of ways, such as superconducting circuits, trapped ions, optical lattices, and linear optics. Computation is performed by operating on the state of these qubits using quantum logic gates. For example, if the qubit is realized as an ion, the quantum logic gate might manipulate the ion's energy state with lasers.

Recall: A metric for classification models. Recall identifies the frequency with which a model correctly classifies the true positive items. It answers the question "How many true positive items were correctly classified"? For example, the percentage of spam messages that were flagged as spam.⁶⁰

Reinforcement Learning: A method of training algorithms to make suitable actions by maximizing rewarded behavior over the course of its actions.⁶¹ This type of learning can take place in simulated environments, such as game-playing, which reduces the need for real-world data.

Reliable AI: An AI system that performs in its intended manner within the intended domain of use.

Responsible AI: An AI system that aligns development and behavior to goals and values. This includes developing and fielding AI technology in a manner that is consistent with democratic values.⁶²

Robotics: A broad field of study including autonomous systems that exist in the physical world, sensing their environment and taking actions to achieve specific goals.⁶³

Robotic Process Automation (RPA): Software to help in the automation of tasks, especially those that are tedious and repetitive.

Robust AI: An AI system that is resilient in real-world settings, such as an object-recognition application that is robust to significant changes in lighting. The phrase also refers to resilience when it comes to adversarial attacks on AI components.

Self-Healing Robots: Robots that use structural materials to self-identify damage and initiate healing on their own, repeatedly.⁶⁴

Self-Replicating Robots: A means of manufacturing, so that fleets of autonomous rovers can extract water and metals from local terrain—say on the moon or Mars—to construct new industrial robots autonomously and continue the self-replication loop.

Self-Supervised Machine Learning: A collection of machine learning techniques that are used to train models or learn embedded representations without reliance on costly labeled data; rather, an approach is to withhold part of each data sample and require the algorithm to learn to predict the missing piece.⁶⁵ Self-supervision has been used to train some of the largest language models built to date by training on large amounts of natural language data.⁶⁶

Semi-Supervised Machine Learning: A process for training an algorithm on a combination of labeled and unlabeled data. Typically, this combination will contain a very small amount of labeled data and a very large amount of unlabeled data. One approach is to use the costly, smaller amount of labeled data to bootstrap a classification model, use that model to generate predicted labels across the larger, unlabeled data, and then use the outcome to retrain/refine the model and iterate until class label assignments stabilize.

Semiconductor Manufacturing Equipment (SME): The tools and equipment required to fabricate semiconductors (e.g., extreme ultraviolet and argon fluoride immersion lithography tools).

Semiconductor Photonics: As it relates to semiconductors, this refers to the use of light, rather than electricity, to transfer information on a chip. This allows for much faster data transfer speeds, resulting in significant performance improvements.

Semiconductors: The silicon-based integrated circuits that drive the operations and functioning of computers and most electronic devices.

Smart Sensors: Devices capable of pre-processing raw data and prioritizing the data to transmit and store, which is especially helpful in degraded or low-bandwidth environments.

Smart Systems: Information technology systems with autonomous functions enabled by AI.

Speech Recognition: The algorithmic process of turning speech signals into text or commands.⁶⁷

Supervised Machine Learning: A process for training algorithms by example. The training data consists of inputs paired with the correct outputs. During training, the algorithm will search for patterns in the data that correlate with the desired outputs and learn to predict the correct output for newly presented input data over iterative training and model updates.

SWaP: Size, weight, and power, typically used in the context of reducing the overall dimensions of a device, increasing its efficiency, and lowering the overall footprint and cost—all contributing factors to viable edge computing.⁶⁸

Symbolic Logic: A tool for creating and reasoning with symbolic representations of objects and propositions based on clearly defined criteria for logical validity.⁶⁹

Synthetic Data Generation: The process of creating artificial data to mimic real sample data sets. It includes methods for data augmentation that automate the process for generating new example data from an existing data set. Synthetic data generation is increasingly utilized to overcome the burden of creating large labeled datasets for testing and at times training deep neural networks.

Technical Baseline: The government's capability to understand underlying technology well enough to make successful acquisition decisions independent of contractors.⁷⁰

TensorFlow: A free and open-source software library for training neural networks and other machine learning architectures, initially developed by Google Brain.

Test and Evaluation, Verification and Validation (TEVV) of AI Systems: A framework for assessing, incorporating methods and metrics to determine that a technology or system satisfactorily meets its design specifications and requirements, and that it is sufficient for its intended use.

Traceability: A characteristic of an AI system enabling a person to understand the technology, development processes, and operational capabilities (e.g., with transparent and auditable methodologies along with documented data sources and design procedures).

Unintended Bias: Ways in which algorithms might perform more poorly than expected (e.g., higher false positives or false negatives), particularly when disparate outcomes are produced (e.g. across categories, classes or groups).

Unsupervised Machine Learning: A process for training a model in which the model learns from the data itself without any data labels. Two common approaches are clustering (in which inherent groupings are discovered) and association (in which rules that describe large portions of the data are discovered).⁷¹

Virtual Reality: A simulated experience in a computer-generated synthetic, artificial world involving immersion, sensory feedback, and interactivity.⁷²

Appendix A - Endnotes

- ¹ See *Adversarial Machine Learning 101*, GitHub/MITRE (last accessed Feb. 18, 2021), <https://github.com/mitre/advm1threatmatrix/blob/master/pages/adversarial-ml-101.md#adversarial-machine-learning-101>; see also Ionut Arghire, *Microsoft, MITRE Release Adversarial Machine Learning Threat Matrix*, Security Week (last accessed Feb. 16, 2021), <https://www.securityweek.com/microsoft-mitre-release-adversarial-machine-learning-threat-matrix>.
- ² GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 169 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ³ See *Glossary*, ISACA (last accessed Feb. 13, 2021), <https://www.isaca.org/resources/glossary>.
- ⁴ Note that for data-driven AI systems, step 2 is expanded and replaced with 2.a) Acquire data to meet the objective, and 2.b) Train the AI system on the data. These two steps are usually repeated, with data acquisition and training continuing until desired performance objectives are attained. For further discussion on the ML lifecycle, see Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, IEEE Computer Society (May 2019), <https://www.microsoft.com/en-us/research/publication/software-engineering-for-machine-learning-a-case-study/>.
- ⁵ The stack of elements listed here is an adaptation from Andrew W. Moore, Martial Hebert, and Shane Shaneman. See Andrew Moore, et al., *The AI Stack: A Blueprint for Developing and Deploying Artificial Intelligence*, Proc. SPIE 10635 (May 4, 2018), <https://doi.org/10.1117/12.2309483>. For a graphical depiction of the AI stack, see *About*, Carnegie Mellon University Artificial Intelligence (last accessed Jan. 1, 2021), <https://ai.cs.cmu.edu/about>.
- ⁶ See **Recital 26 EU General Data Protection Regulation (EU-GDPR)**, PrivazyPlan (last accessed Feb. 17, 2021), <https://www.privacy-regulation.eu/en/recital-26-GDPR.htm>.
- ⁷ Vinton G. Cerf, *APIs, Standards, and Enabling Infrastructure*, Communications of the ACM, Vol. 62 No. 5, at 5 (May 2019), <https://m-cacm.acm.org/magazines/2019/5/236425-apis-standards-and-enabling-infrastructure/fulltext?mobile=true>.
- ⁸ GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 169 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ⁹ See *Augmented Reality*, Google (last accessed Feb. 13, 2021), <https://arvr.google.com/ar/>.
- ¹⁰ See *Authorization to Operate*, NIST Computer Security Resource Center (last accessed Feb. 13, 2021), https://csrc.nist.gov/glossary/term/authorization_to_operate.
- ¹¹ See *Biosensors*, Nature (last accessed Feb. 13, 2021), <https://www.nature.com/subjects/biosensors>.
- ¹² Maria Korolov, *What Is Biometrics? 10 Physical and Behavioral Identifiers That Can Be Used for Authentication*, CSO (Feb. 12, 2019), <https://www.csoonline.com/article/3339565/what-is-biometrics-and-why-collecting-biometric-data-is-risky.html>.
- ¹³ See *Understanding Cloud Computing*, Red Hat (last accessed Feb. 13, 2021), <https://www.redhat.com/en/topics/cloud>.
- ¹⁴ See *What Is Cloud Infrastructure?*, Red Hat (last accessed Feb. 13, 2021), <https://www.redhat.com/en/topics/cloud-computing/what-is-cloud-infrastructure>.
- ¹⁵ See Matt Turek, *Machine Common Sense (MCS)*, DARPA (last accessed Feb. 13, 2021), <https://www.darpa.mil/program/machine-common-sense>.
- ¹⁶ See Alfred V. Aho, *Ubiquity Symposium: Computational and Computational Thinking*, ACM (January 2011), <https://ubiquity.acm.org/article.cfm?id=1922682>.
- ¹⁷ See *Computer Vision: What It Is and Why It Matters*, SAS (last accessed Feb. 13, 2021), https://sas.com/en_in/insights/analytics/computer-vision.html.
- ¹⁸ GAO-20-590G, *Agile Assessment Guide*, U.S. Government Accountability Office at 171 (Sept. 2020), <https://www.gao.gov/assets/710/709711.pdf>.
- ¹⁹ *Id.* at 172.

²⁰ See Thor Olavsrud, *What Is Data Architecture? A Framework for Managing Data*, CIO (Nov. 4, 2020), <https://www.cio.com/article/3588155/what-is-data-architecture-a-framework-for-managing-data.html>.

²¹ *Digital Strategy 2020-2024*, USAID at 48 (June 2020), https://www.usaid.gov/sites/default/files/documents/15396/USAID_Digital_Strategy.pdf.

²² *Id.*

²³ See Jake Frankenfield, *De-Anonymization*, Investopedia (Dec. 27, 2020), <https://www.investopedia.com/terms/d/deanonymization.asp#:~:text=De%2Danonymization%20is%20a%20technique,person%2C%20group%2C%20or%20transaction>.

²⁴ *Interim Report*, NSCAI at 9 (Nov. 2019), https://www.nscai.gov/wp-content/uploads/2021/01/NSCAI-Interim-Report-for-Congress_201911.pdf.

²⁵ See Ian Goodfellow, et al., *Deep Learning*, MIT Press, (2016), <https://www.deeplearningbook.org/>.

²⁶ *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD*, MITRE, at 9-25 (Jan. 2017), <https://fas.org/irp/agency/dod/jason/ai-dod.pdf>.

²⁷ See *Understanding the Differences Between Agile and DevSecOps—From a Business Perspective*, General Services Administration (last accessed Feb. 13, 2021), https://tech.gsa.gov/guides/understanding_differences_agile_devsecops/.

²⁸ Kobbi Nissim, et al., *Differential Privacy: A Primer for a Non-technical Audience*, Working Group of the Privacy Tools for Sharing Research Data Project, Harvard University (Feb. 14, 2018), https://privacytools.seas.harvard.edu/files/privacytools/files/pedagogical-document-dp_new.pdf.

²⁹ *Digital Strategy 2020-2024*, USAID at 4 (June 2020), https://www.usaid.gov/sites/default/files/documents/15396/USAID_Digital_Strategy.pdf.

³⁰ *Id.* at 49.

³¹ Maarten van Steen & Andrew Tanenbaum, *Distributed Systems* (3rd ed.), distributed-systems.net (2017), <https://www.distributed-systems.net/index.php/books/ds3/>.

³² See Eric Hamilton, *What Is Edge Computing: The Network Edge Explained*, Cloudwards (Dec. 27, 2018), <https://www.cloudwards.net/what-is-edge-computing>.

³³ Peter Jackson, *Introduction to Expert Systems* (3rd ed.), Addison Wesley at 2 (1998).

³⁴ For further discussion see P. Jonathon Phillips, et al., *Four Principles of Explainable Artificial Intelligence*, NIST (Aug. 2020), <https://www.nist.gov/system/files/documents/2020/08/17/NIST%20Explainable%20AI%20Draft%20NISTIR8312%20%281%29.pdf>.

³⁵ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

³⁶ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

³⁷ Ian Goodfellow, et al., *Generative Adversarial Nets*, Neural Information Processing Systems (2014), <https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>.

³⁸ *Fiscal Year 2019: Stockpile Stewardship and Management Plan—Biennial Plan Summary, Report to Congress*, U.S. Department of Energy at 3-7 (Oct. 2018), <https://www.energy.gov/sites/prod/files/2018/10/f57/FY2019%20SSMP.pdf>.

³⁹ See *Introduction*, Homomorphic Encryption Standardization (last accessed Feb. 13, 2021), <https://homomorphicencryption.org/introduction/>.

⁴⁰ Eric Horvitz, *Principles of Mixed-Initiative User Interfaces*, CHI '99: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems at 159-166 (May 1999), <https://dl.acm.org/doi/pdf/10.1145/302979.303030>.

Appendix A - Endnotes

- ⁴¹ Bryan Wilder, et al., *Learning to Complement Humans*, Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence (IJCAI-20) at 1526-1533 (Jan. 2021), <https://www.ijcai.org/Proceedings/2020/0212.pdf>.
- ⁴² Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI '19: Proceedings of the CHI Conference on Human Factors in Computing Systems, at 1-13 (May 2019), <https://dl.acm.org/doi/pdf/10.1145/3290605.3300233>.
- ⁴³ Catherine Theohary, *Defense Primer: Information Operations*, Congressional Research Service (Dec. 15, 2020), <https://fas.org/sgp/crs/natsec/IF10771.pdf>.
- ⁴⁴ See interactive *ITU Terms and Definitions*, United Nations International Telecommunication Union (last accessed Feb. 15, 2021), <https://www.itu.int/net/ITU-R/asp/terminology-definition.asp?lang=en&rlink={42AA741E-A0A7-48C4-905B-AAAFDA29E5F2}>.
- ⁴⁵ See *Intelligent Sensors*, MDPI Sensors (last accessed Feb. 13, 2021), https://www.mdpi.com/journal/sensors/sections/Intelligent_Sensors.
- ⁴⁶ See Richard Gall, *Machine Learning vs Interpretability: Two Concepts That Could Help Restore Trust in AI*, KDnuggets (Dec. 2018), <https://www.kdnuggets.com/2018/12/machine-learning-explainability-interpretability-ai.html>.
- ⁴⁷ A. Sivagnana Ganesan & T. Chithralekha, *A Survey on Survey of Migration of Legacy Systems*, ICIA-16: Proceedings of the International Conference on Informatics and Analytics at 1-10 (Aug. 2016), <https://dl.acm.org/doi/10.1145/2980258.2980409>.
- ⁴⁸ Thomas M. Mitchell, *Machine Learning*, McGraw-Hill (1997).
- ⁴⁹ See *2021 Technology Spotlight: The Emergence of MLOps*, Booz Allen Hamilton (2021), https://www.boozallen.com/content/dam/boozallen_site/dig/pdf/white_paper/the-emergence-of-mlops.pdf.
- ⁵⁰ Peter Kairouz, et al., *Advances and Open Problems in Federated Learning*, arXiv (Dec. 10, 2019), <https://arxiv.org/pdf/1912.04977.pdf>.
- ⁵¹ See Vale Tolpegin et al., *Data Poisoning Attacks Against Federated Learning Systems*, ArXiv (Aug. 11, 2020), <https://arxiv.org/abs/2007.08432>; Arjun Nitin Bhagoji, et al., *Analyzing Federated Learning Through an Adversarial Lens*, arXiv (Nov. 25, 2019), <https://arxiv.org/abs/1811.12470>.
- ⁵² See *Beyond Today's AI: New Algorithmic Approaches Emulate the Human Brain's Interactions with the World*, Intel (last accessed Feb. 13, 2021), <https://www.intel.com/content/www/us/en/research/neuromorphic-computing.html>.
- ⁵³ Ramesh Jain, et al., *Machine Vision*, McGraw-Hill at 459 (1995), https://www.cse.usf.edu/~r1k/MachineVisionBook/MachineVision.files/MachineVision_Chapter15.pdf.
- ⁵⁴ Adam Santoro, et al., *One-Shot Learning with Memory-Augmented Neural Networks*, arXiv (May 19, 2016), <https://arxiv.org/pdf/1605.06065.pdf>.
- ⁵⁵ See *About Workshop, Open Knowledge Network* at National Institutes of Health, Subcommittee on Networking & Information Technology Research & Development, Big Data Interagency Working Group, (Oct. 4-5, 2017), https://www.nitrd.gov/nitrdgroups/index.php?title=Open_Knowledge_Network.
- ⁵⁶ Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer at 1 (2006), <https://www.microsoft.com/en-us/research/uploads/prod/2006/01/Bishop-Pattern-Recognition-and-Machine-Learning-2006.pdf>.

- ⁵⁷ See Nate Lord, *What Is Polymorphic Malware? A Definition and Best Practices for Defending Against Polymorphic Malware*, Digital Guardian (July 17, 2020), <https://digitalguardian.com/blog/what-polymorphic-malware-definition-and-best-practices-defending-against-polymorphic-malware#:~:text=Definition%20of%20Polymorphic%20Malware.bots%2C%20trojans%2C%20or%20keyloggers.>
- ⁵⁸ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.
- ⁵⁹ Trevor Hastie, et al., *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.), Springer at 9-11 (Jan. 13, 2017), https://web.stanford.edu/~hastie/ElemStatLearn/printings/ESLII_print12_toc.pdf.
- ⁶⁰ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2021), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.
- ⁶¹ Richard S. Sutton & Andrew G. Barto, *Reinforcement Learning: An Introduction* (2nd ed.), MIT Press (2018).
- ⁶² *Key Considerations for Responsible Development and Fielding of Artificial Intelligence*, NSCAI (July 22, 2020), <https://www.nscai.gov/previous-reports/>.
- ⁶³ See Erico Guizzo, *What Is a Robot?*, IEEE (May 28, 2020), <https://robots.ieee.org/learn/what-is-a-robot/>.
- ⁶⁴ See Evan Ackerman, *Soft Self-Healing Materials for Robots That Cannot Be Destroyed*, IEEE (Sept. 5, 2019), <https://spectrum.ieee.org/automaton/robotics/robotics-hardware/soft-selfhealing-materials-for-robots-that-cannot-be-destroyed>.
- ⁶⁵ See Andrew Zisserman, *Self-Supervised Learning*, Google DeepMind (last accessed Feb. 17, 2021), <https://project.inria.fr/paiss/files/2018/07/zisserman-self-supervised.pdf>.
- ⁶⁶ Jacob Devlin, et al., *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding*, arXiv (May 24, 2019), <https://arxiv.org/pdf/1810.04805.pdf>.
- ⁶⁷ Jianliang Meng, et al., *Overview of the Speech Recognition Technology*, 2012 Fourth International Conference on Computational and Information Sciences at 199-202 (2012), <https://ieeexplore.ieee.org/document/6300437/>.
- ⁶⁸ See *What Is Low-SWaP?*, REDCOM (last accessed Feb. 13, 2021), <https://www.redcom.com/what-is-low-swap-size-weight-and-power/>.
- ⁶⁹ Tony Roy, *Symbolic Logic: An Accessible Introduction to Serious Mathematical Logic* at 2-3 (Feb. 8, 2021), <https://tonyroypphilosophy.net/symbolic-logic/>.
- ⁷⁰ William LaPlante, *Owning the Technical Baseline*, Defense AT&L at 18-20, (July-Aug. 2015), <https://apps.dtic.mil/dtic/tr/fulltext/u2/1016084.pdf>.
- ⁷¹ See Jason Brownlee, *Supervised and Unsupervised Machine Learning Algorithms*, in *Machine Learning Algorithms*, Machine Learning Mastery (Aug. 20, 2020), <https://machinelearningmastery.com/supervised-and-unsupervised-machine-learning-algorithms/>.
- ⁷² See *What Is Virtual Reality?* Virtual Reality Society (last accessed Feb. 13, 2021), <https://www.vrs.org.uk/virtual-reality/what-is-virtual-reality.html>.

Appendix B: Acronyms Found in This Report

(alphabetical order)

A	
AAAI	Association for the Advancement of Artificial Intelligence
AaaS	applications as a service
AAMAS	Autonomous Agents and Multiagent Systems
AISC	AI Strategic Challenge
AAF	Adaptive Acquisition Framework
AAL	Army Applications Laboratory
ABMS	Advanced Battle Management System
ADL	Advanced Distributed Learning
AFC	Army Futures Command
AFOSR	Air Force Office of Scientific Research
AFRL	Air Force Research Lab
AGI	artificial general intelligence
ACM SIGKDD	Association for Computing Machinery's Special Interest Group on Knowledge Discovery and Data Mining
AI	artificial intelligence
AI CoE	AI Center of Excellence
AIM	Augmenting Intelligence using Machines
AIPfd	AI Partnership for Defense

Amii	Alberta Machine Intelligence Institute
ANPRM	Advance Notice of Proposed Rulemaking
API	Application Programming Interface
ArF	Argon fluoride
ARO	Army Research Office
ARPA	Academic Research Protection Act
ASC	Alternative Simplified Credit
ASIC	application-specific integrated circuit
ATO	Authorization (or Authority) to Operate
AVC	Bureau of Arms Control, Verification and Compliance

B

BA	Budget Activity
BARDA	Biomedical Advanced Research and Development Authority
BioMADE	Bioindustrial Manufacturing and Design Ecosystem
BIRD	Binational Industrial Research & Development Foundation
BIS	Bureau of Industry and Security
BSF	Binational Science Foundation

C

C2	command and control
C&ET	critical and emerging technologies
CBP	U.S. Customs and Border Protection
CBRS	Citizens Broadband Radio Service
CCMD	combatant command
CCP	Chinese Communist Party

CCW	Convention on Certain Conventional Weapons
CD	cardiovascular disease
CDC	Centers for Disease Control and Prevention
CDO	chief data officer
CFIUS	Committee on Foreign Investment in the United States
CFR	Code of Federal Regulations
CHIPS	Creating Helpful Incentives to Produce Semiconductors
CIA	Central Intelligence Agency
CIO	chief information officer
CISA	Cybersecurity and Infrastructure Security Agency
CMI	Component Mission Initiative
CNIPA	China National Intellectual Property Administration
COE	Center of Excellence
CONOPS	concept(s) of operations(s)
COTS	commercial off-the-shelf
COVID-19	coronavirus disease 2019
CReATE	Coding Repository and Transformation Environment
CRISPR	clustered regularly interspaced short palindromic repeats
CRCL	civil rights and civil liberties
CS	computer science
CSC	U.S. Cyberspace Solarium Commission
CSET Bureau	Bureau of Cyberspace Security and Emerging Technologies
CSIS	Center for Strategic and International Studies
CSTD	Comprehensive Science and Technology Dialogue
CTO	chief technology officer

D

DA	Decision Authority
DAC	Development Assistance Committee
DARPA	Defense Advanced Research Projects Agency
DDI	Bureau for Development, Democracy, and Innovation at USAID
DEXCOM	Deputies Executive Committee
DFC	U.S. International Development Finance Corporation
DFFT	data free flow with trust
DFI	development finance institution
DIA	Defense Intelligence Agency
DIB	Defense Innovation Board
DIU	Defense Innovation Unit
DHS	Department of Homeland Security
D/MR	Deputy Secretary of State for Management and Resources
DNA	deoxyribonucleic acid
DNI	Director of National Intelligence
DoD	Department of Defense
DoDD	Department of Defense Directive
DoE	Department of Energy
DOI	Department of the Interior
DOJ	Department of Justice
DOT	Department of Transportation
DOT&E	Director, Operational Test and Evaluation
DOTMLPF-P	doctrine, organization, training, materiel, leadership and education, personnel, facilities, and policy
DPC	Domestic Policy Council

DRL	Bureau of Democracy, Human Rights, and Labor
E	
EB	Bureau of Economic and Business Affairs
ECRA	Export Control Reform Act of 2018
EDT	emerging and disruptive technology
E.O.	Executive Order
EOP	Executive Office of the President
ERI	Electronics Resurgence Initiative
ESA	European Space Agency
ETC	Emerging Technology Coalition
ETTAC	Emerging Technology Technical Advisory Committee
EU	European Union
EUV	extreme ultraviolet
EXIM	Export-Import Bank of the United States
F	
FAIR	Facebook AI Research
FAR	Federal Acquisition Regulation
FARA	Foreign Agents Registration Act
FBI	Federal Bureau of Investigation
FCC	Federal Communications Commission
FCEN	Financial Crimes Enforcement Network
FDA	U.S. Food and Drug Administration
FFRDC	Federally Funded Research and Development Center
FIRRMA	Foreign Investment Risk Review Modernization Act of 2018

FISMA	Federal Information Security Modernization Act
FPGA	field-programmable gate array
FSI	Foreign Service Institute
FTQC	fault-tolerant quantum computer
FWCI	field-weighted citation impact
FY	fiscal year
FYDP	Future Years Defense Plan
G	
G20	Group of 20
GAN	generative adversarial network
GAO	U.S. Government Accountability Office
GDP	gross domestic product
GEC	Global Engagement Center at Department of State
GGE	Group of Governmental Experts
GIST	Global Innovation through Science and Technology
GPAI	Global Partnership on Artificial Intelligence
GPS	Global Positioning System
GPT-3	Generative Pre-trained Transformer 3
GPU	graphics processing unit
GSA	U.S. General Services Administration
H	
HPC	high-performance computing
HHMI	Howard Hughes Medical Institute
HHS	Health and Human Services

HR	human resources
HQE	highly qualified expert
HSI	human-system interactions
HUMINT	human intelligence
I	
I&W	indication(s) and warning(s)
IARPA	Intelligence Advanced Research Projects Activity
IC	U.S. Intelligence Community
IC ITE	Intelligence Community Information Technology Environment
ICRC	International Committee of the Red Cross
ICT	information and communications technology
IDDI	International Digital Democracy Initiative
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronic Engineers
IER	International Entrepreneur Rule
IFBHR	Internet Freedom and Business & Human Rights Section
IFI	international financial institution
IHL	International Humanitarian Law
IMINT	imagery intelligence
INL	Bureau of International Narcotics and Law Enforcement Affairs
IoT	internet of things
IP	intellectual property
IPEC	U.S. Intellectual Property Enforcement Coordinator
IPA	Intergovernmental Personnel Act
IPHE	International Partnership for Hydrogen and Fuel Cells in the Economy

ISAC	Information Sharing and Analysis Center
ISIS	Islamic State of Iraq and Syria
ISN	Bureau of International Security and Nonproliferation
ISO	International Organization for Standardization
ISR	intelligence, surveillance, and reconnaissance
ISTS	International Science and Technology Strategy
IT	information technology
ITF-CCAD	International Task Force to Counter and Compete Against Disinformation
IT SRMC	IT Modernization Senior Risk Management Council
ITU	International Telecommunication Union
ITU-T	International Telecommunication Union–Telecommunication Standardization Sector
IUSSTF	Indo-U.S. Science and Technology Forum

J

JAIC	Joint Artificial Intelligence Center
JCF	Joint Common Foundation
JCIDS	Joint Capabilities Integration and Development System
JIATF	Joint Interagency Task Force
JROC	Joint Requirements Oversight Council
JSP	Joint Strategic Plan
JWAC	Joint Warfare Analysis Center

K

K-12	kindergarten to 12th grade
------	----------------------------

L

LAWS	lethal autonomous weapon systems
LKIE	learning, knowledge, and information exchange
LOAC	Law of Armed Conflict

M

M&A	mergers and acquisitions
M&S	modeling and simulation
MAIEI	Montreal AI Ethics Institute
MAIRI	Multilateral AI Research Institute
MASINT	Measurement and signature intelligence
MCC	Millennium Challenge Corporation
MDA	Milestone Decision Authorities
MDAP	Major Defense Acquisition Program
MediFOR	Media Forensics
MEMT	Multi-Engine Machine Translation
Mila	Montreal Institute for Learning Algorithms
MIT	Massachusetts Institute of Technology
MITE	Malign Information Threat Executive
ML	machine learning

N

NAIRR	National Artificial Intelligence Research Resource
NASA	National Aeronautics and Space Administration
NATO	North Atlantic Treaty Organization
NBA	National Basketball Association

NCEAI	National Council for Expanding American Innovation
NCPS	National Cybersecurity Protection System
NCSC	National Counterintelligence and Security Center
NCTC	National Counterterrorism Center
NDAA	National Defense Authorization Act
NDEA	National Defense Education Act
NDS	National Defense Strategy
NEA	Nuclear Energy Agency
NEC	National Economic Council
NIFA	National Institute of Food and Agriculture
NIH	National Institutes of Health
NIJ	National Institute of Justice
NIS	National Intelligence Strategy
NISQ	noisy intermediate-scale quantum
NIST	National Institute of Standards and Technology
NITRD	Networking and Information Technology Research and Development
NLP	natural language processing
NLU	natural language understanding
NOAA	National Oceanic and Atmospheric Administration
NQCO	National Quantum Coordination Office
NQI	National Quantum Initiative
NRDC	National Reserve Digital Corps
NSF	National Science Foundation
NSTC	National Science and Technology Council
NSA	National Security Agency

NSC	National Security Council
NSF	National Science Foundation
NSIB	National Security Innovation Base
NSIN	National Security Innovation Network
NSS	National Security Strategy
NTF	National Technology Foundation
NTS	National Technology Strategy
NTIA	National Telecommunications and Information Administration
NVLAP	National Voluntary Laboratory Accreditation Program
O	
ODNI	Office of the Director of National Intelligence
OECD	Organisation for Economic Co-operation and Development
OES	Bureau of Oceans and International Environmental and Scientific Affairs
OISE	Office of International Science and Engineering
OMB	Office of Management and Budget
ONR	Office of Naval Research
OPM	U.S. Office of Personnel Management
ORSA	operational research and systems analysis
OSD	Office of the Secretary of Defense
OSINT	open-source intelligence
OSTP	Office of Science and Technology Policy
OTA	Other Transaction Authority
OUSD (A&S)	Office of the Under Secretary of Defense for Acquisition and Sustainment
OUSD (I&S)	Office of the Under Secretary of Defense for Intelligence and Security
OUSD (R&E)	Office of the Under Secretary of Defense for Research and Engineering

OZ	Opportunity Zone
P	
PaaS	platforms as a service
PAL	Permissive Action Link
PCAST	President's Council of Advisors on Science and Technology
P/CLR	Privacy, Civil Liberties, and Civil Rights
P/CRCL	Privacy, Civil Rights, and Civil Liberties
PCLOB	Privacy and Civil Liberties Oversight Board
PCT	Patent Cooperation Treaty
PDDNI	Principal Deputy Director of National Security
PE	program element
PED	processing, exploitation, and dissemination
PGNN	physics-guided neural network
PhD	doctoral graduate
PIA	Privacy Impact Assessment
PII	personally identifiable information
PLA	People's Liberation Army
PM	Bureau of Political-Military Affairs
PM	program manager
PoR	Program of Record
PPBE	Planning, Programming, Budget, and Execution
PPML	privacy-preserving machine learning

Q

QED-C	Quantum Economic Development Consortium
-------	---

QIS	Quantum Information Science
-----	-----------------------------

QPU	quantum processing unit
-----	-------------------------

R

R&D	research and development
-----	--------------------------

RAI	responsible AI
-----	----------------

RDT&E	research, development, test, and evaluation
-------	---

REN-ISAC	Research and Education Networks Information and Sharing Analysis Center
----------	---

RL	reinforcement learning
----	------------------------

RMF	Risk Management Framework
-----	---------------------------

RPA	robotic process automation
-----	----------------------------

S

S&E	science and engineering
-----	-------------------------

SBIR	Small Business Innovation Research Program
------	--

S/CCI	Office of the Coordinator for Cyber Issues
-------	--

SDK	software development kit
-----	--------------------------

SDO	standards developing organization
-----	-----------------------------------

SemaFor	semantic forensics
---------	--------------------

SEP	“standard essential” patents
-----	------------------------------

SFS	scholarship for service
-----	-------------------------

SGE	Special Government Employee
-----	-----------------------------

SIAC	Strategic Intelligence Analysis Cell
------	--------------------------------------

SIGINT	signals intelligence
--------	----------------------

SMART	Science, Mathematics, and Research for Transformation
SME	semiconductor manufacturing equipment
SMIC	Semiconductor Manufacturing International Corporation
SORN	System of Records Notice
SSD	Strategic Security Dialogue
S&T	science and technology
STAS	Office of the Science and Technology Adviser to the Secretary of State
State/Q	Under Secretary of State for Science, Research and Technology
STEM	science, technology, engineering, and mathematics
STTR	Small Business Technology Transfer Program
SWaP	size, weight, and power

T

TCC	Technology Competitiveness Council
T&E	test(ing) and evaluation
TET	Technology Engagement Team at Department of State
TEVV	test(ing) and evaluation, verification and validation
TRC	Technology Research Center
TSMC	Taiwan Semiconductor Manufacturing Corporation
TTCP	Technical Cooperation Program

U

UARC	University Affiliated Research Center
U.K.	United Kingdom
UN	United Nations
U.S.	United States

U.S.C.	United States Code
USAF	U.S. Air Force
USAID	U.S. Agency for International Development
USASOC	U.S. Army Special Operations Command
USCIS	U.S. Citizenship and Immigration Services
USD(A&S)	Under Secretary of Defense for Acquisition and Sustainment
USD(R&E)	Under Secretary of Defense for Research and Engineering
USDA	U.S. Department of Agriculture
USDSA	U.S. Digital Service Academy
USERRA	Uniformed Services Employment and Reemployment Rights Act
USISTEF	United States–India Science & Technology Endowment Fund
USPTO	U.S. Patent and Trademark Office
USTDA	U.S. Trade and Development Agency
USTR	Office of the U.S. Trade Representative

V

VA U.S. Department of Veterans Affairs

VCJCS Vice Chairman of the Joint Chiefs of Staff

W

WH White House

WIPO World Intellectual Property Organization

WLIF Warfighting Lab Incentive Fund

Numbers

3D three-dimensional

3SIIF Three Seas Initiative Investment Fund

5G fifth-generation standard for broadband cellular networks

Appendix C: Key Considerations for the Responsible Development and Fielding of Artificial Intelligence (Abridged)

Prefatory Note:

The paradigm and recommended practices described here stem from the Commission's line of effort dedicated to Ethics and Responsible Artificial Intelligence (AI). The Commission has recommended that heads of departments and agencies critical to national security (at a minimum, the Department of Defense, Intelligence Community, Department of Homeland Security, Federal Bureau of Investigation, Department of Energy, Department of State, and Department of Health and Human Services) should implement the Key Considerations as a paradigm for the responsible development and fielding of AI systems. This includes developing processes and programs aimed at adopting the paradigm's recommended practices, monitoring their implementation, and continually refining them as best practices evolve.

This approach would set the foundation for an intentional, government-wide, coordinated effort to incorporate recommended practices into current processes for AI development and fielding. However, our overarching aim is to allow agencies to continue to have the flexibility to craft policies and processes according to their specific needs. The Commission is mindful of the required flexibility that an agency needs when conducting the risk assessment and management of an AI system, as these tasks will largely depend on the context of the AI system.

This recommendation, along with a set of recommended considerations and practices, was made originally in July 2020. Here we present a revised and updated version as part of the Commission's Final Report. Many of the points made here are also reflected in Chapter 7 of the report.

The content herein is an abridged version of the content included in the extended version, which will be featured on NSCAI's website in March 2021 at www.nsc.ai.gov. In the more comprehensive document, we provide additional details and references for technical implementers.

Introduction

The Commission acknowledges the efforts undertaken to date to establish ethics guidelines for AI systems.¹ While some national security agencies have adopted,² or are in the process of adopting, AI principles,³ other agencies have not provided such guidance. In cases where principles are offered, it can be difficult to translate the high-level concepts into concrete actions. In addition, agencies would benefit from the establishment of greater consistency in policies to further the responsible development and fielding of AI technologies across government.

This Commission has identified five broad categories of challenges and made recommendations for both responsibly developing and fielding AI systems. These recommendations include immediate actions and future work the U.S. government should undertake to help establish best practices to overcome these challenges. Collectively, they form a paradigm for aligning AI system development and AI system behavior to goals and values. The first section, *Aligning Systems and Uses with American Values and the Rule of Law*, provides guidance specific to implementing systems that abide by American values, most of which are shared by democratic nations. The section also covers aligning the run-time behavior of systems to the related, more technical encodings of objectives, utilities, and trade-offs. The four following sections (on *Engineering Practices*, *System Performance*, *Human-AI Interaction*, and *Accountability & Governance*) serve in support of core American values and further outline practices needed to develop and field AI systems that are understandable, reliable, robust, and trustworthy.

Recommended practices span multiple phases of the AI lifecycle and establish a baseline for the responsible development and fielding of AI technologies. The Commission uses “development” to refer to “designing, building, and testing during development and prior to deployment” and “fielding” to refer to “deployment, monitoring, and sustainment.”

The Commission recommends that heads of departments and agencies implement the Key Considerations as a paradigm for the responsible development and fielding of AI systems. This includes developing policies and processes to adopt the paradigm's recommended practices, monitor their implementation, and continually refine them as best practices evolve. These recommended practices should apply both to systems that are developed by departments and agencies as well as to those that are acquired. Systems acquired (whether commercial off-the-shelf systems or through contractors) should be subjected to the same rigorous standards and recommended practices in the acquisitions and acceptance processes. As such, the government organization overseeing the bidding

process should require that vendors articulate how their practices align with the Key Considerations' recommended practices in their proposals, submissions, and bids.

In each of the five sections that follow, we first provide a conceptual overview of the scope and importance of the topic. We then illustrate examples of a current challenge relevant to national security departments that underscores the need to adopt recommended practices in this area. Then, we provide a list of recommended practices that agencies should adopt, acknowledging research, industry tools, and exemplary models within government that could support agencies in the adoption of recommended practices. Finally, in areas where best practices do not exist or are especially challenging to implement, we note the need for future work as a priority; this includes, for example, R&D and standards development. We also identify potential areas in which collaboration with allies and partners would be beneficial for interoperability and trust and note that the Key Considerations can inform potential future efforts to discuss military uses of AI with strategic competitors.

I. Aligning Systems and Uses with American Values and the Rule of Law

(1) Overview

Our values guide our decisions and our assessment of their outcomes. Our values shape our policies, our sensitivities, and how we balance trade-offs among competing interests. America's values, and commitment to upholding them, are reflected in the U.S. Constitution and U.S. laws, regulations, policies, and processes.

One of the seven principles we set forth in the Commission's Interim Report (November 2019) is the following:

The American way of AI must reflect American values—including having the rule of law at its core. For federal law enforcement agencies conducting national security investigations in the United States, that means using AI in ways that are consistent with constitutional principles of due process, individual privacy, equal protection, and non-discrimination. For American diplomacy, that means standing firm against uses of AI by authoritarian governments to repress individual freedom or violate the human rights of their citizens. And for the U.S. military, that means finding ways for AI to enhance its ability to uphold the laws of war and ensuring that current frameworks adequately cover AI.

Values established in the U.S. Constitution, and further operationalized in legislation, include freedoms of speech and assembly as well as the rights to due process, inclusion, fairness, non-discrimination (including equal protection), and privacy (including protection from unwarranted government interference in one's private affairs). These values are codified in the U.S. Constitution and the U.S. Code.⁴ International treaties that the United States has ratified also demonstrate our values by affirming our commitments to human rights and human dignity.⁵ Within America's national security departments, our commitment to

protecting and upholding privacy and civil liberties is further embedded in the policies and programs of the Intelligence Community (IC),⁶ the Department of Homeland Security,⁷ the Department of Defense (DoD),⁸ and oversight entities (e.g., the Privacy and Civil Liberties Oversight Board).⁹ In the military context, core values such as distinction and proportionality are embodied in the nation's commitment to, and the DoD's policies to uphold, the Uniform Code of Military Justice and the Law of Armed Conflict (LOAC).¹⁰

Other values are reflected in treaties, rules, and policies, such as the Convention Against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment¹¹; the DoD's Rules of Engagement¹²; and the DoD's directive concerning autonomy in weapon systems.¹³ While not an exhaustive list of U.S. values, the paradigm of considerations and recommended practices for AI that we introduce resonates with these values, as they have been acknowledged as critical by the U.S. government and national security departments and agencies. Further, many of these values are common to America's like-minded partners, who share a commitment to democracy, human dignity, and human rights.

Our values demand that the development and fielding of AI respect these foundational values and that they enable human empowerment as well as accountability. They require that the operation of AI systems and components be compliant with our laws and international legal commitments and with our departmental policies. In short, American values must inform the way we develop and field AI systems and the way our AI systems behave in the world.

(2) Examples of Current Challenges

Machine learning (ML) techniques can assist DoD agencies with large-scale data analyses to support and enhance decision-making about personnel. As an example, the Proposed New Disability Construct (PNDC) seeks to leverage data analyses to identify service members on the verge of ineligibility for deployment due to concerns with their readiness. Other potential analyses, including factors that lead to success or failure in promotion, can support personnel evaluations. Caution and proven practices are needed, however, to avoid pitfalls in fairness and inclusiveness, several of which have been highlighted in high-profile challenges in areas like criminal justice, recruiting and hiring, and face recognition.¹⁴ Attention should be paid to challenges with decision support systems like PNDC to avoid harmful disparate impact.¹⁵ Likewise, factors weighed in performance evaluations and promotions must be carefully considered to avoid inadvertently reinforcing existing biases through ML-assisted decisions.¹⁶

(3) Recommendations for Adoption

A. *Developing uses and building systems that behave in accordance with American values and the rule of law.* To implement core American values, it is important to:

1. *Employ technologies and operational policies that align with privacy preservation, fairness, inclusion, human rights, and the law of armed conflict (LOAC).* Technologies and policies throughout the AI lifecycle should support achieving these goals. They should ensure that AI uses and systems are consistent with these values and mitigate the risk that AI system uses/outcomes will violate these values.

- An explicit analysis of outcomes that would violate these values should be performed. Policy should prohibit disallowed outcomes that would violate the values above. During system development, analysis of system-specific disallowed outcomes should be performed.¹⁷ As the technology advances, applications evolve, and our understanding of the implications of use grows, these policies should periodically be refreshed.

B. *Representing objectives and trade-offs.* Another important practice for aligning AI systems with values is to consider values as (1) embodied in choices about engineering trade-offs and (2) explicitly represented in the goals and utility functions of an AI system.¹⁸ Recommended practices for representing objectives and trade-offs include the following:

1. *Consider and document value considerations in AI systems by specifying how trade-offs with accuracy are handled.* This includes documenting the choices made when selecting operating thresholds that have implications for performance, such as the ratio of true positive and false positive rates or the precision (how many selected items are relevant?) versus recall (how many relevant items are selected?). For example, consider a system designed to recommend if a person entering the U.S. should be pulled aside for more detailed inspection and interview. Precision refers to how many of the people selected for additional processing are valid security concerns; recall refers to how many valid security concerns are flagged for added processing. The trade-off is between allowing a valid security concern to slip past review and detaining persons who are not a security concern. Setting thresholds to increase precision (i.e., reduce the number of persons detained needlessly) will drive down recall (i.e., detain fewer valid security concerns).

2. *Consider and document value considerations in AI systems that rely on representations of objective or utility functions,* especially when assigning weighting that captures the importance of different goals for the system. As an illustration of multiple goals and value weights, consider shopping for a new car. A buyer may identify factors that are important in the decision, such as gas mileage, safety, reliability, and performance. These clearly interact in some cases—for example, gas mileage and performance are likely in tension, and safety is likely correlated partly with vehicle size, which is likely in

tension with gas mileage. When reviewing a set of new cars, the best pick for a buyer will depend on the priorities placed on these factors.

3. *Conduct documentation, reviews, and set limits that reflect disallowed outcomes (through constraints on allowed performance) to ensure compliance with values.*

(4) Recommendations for Future Action

Future R&D. R&D is needed to advance capabilities for preserving and ensuring that developed or acquired AI systems will act in accordance with American values and the rule of law. For instance, the Commission notes the need for R&D to assure that the personal privacy of individuals is protected in the acquisition and use of data for AI system development. This includes advancing ethical practices with the use of personal data, including disclosure and consent about data collection and use models (including uses of data to build base models that are later retrained and fine-tuned for specific tasks), the use of anonymity techniques and privacy-preserving technologies, and uses of related technologies such as multiparty computation (to allow collaboration on the pooling of data from multiple organizations without sharing data sets). Additionally, we need to understand the compatibility of data usage policies and privacy-preserving approaches with regulatory approaches such as the European Union's General Data Protection Regulation (GDPR).

II. Engineering Practices

(1) Overview

The government and its partners (including vendors), should adopt recommended practices for creating and maintaining trustworthy and robust AI systems that are *auditable* (able to be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations¹⁹); *traceable* (to understand the technology, development processes, and operational methods applicable to AI capabilities, for example with transparent and auditable methodologies, data sources, and design procedure and documentation²⁰); *interpretable* (to understand the value and accuracy of system output²¹); and *reliable* (to perform in the intended manner within the intended domain of use²²). There are no broadly directed best practices or standards to guide organizations in the building of AI systems that are consistent with designated AI principles, but potential approaches, minimal standards, and engineering proven practices are available.²³

Additionally, several properties of the engineering methods and models used in ML (e.g., data-centric methods) are associated with weaknesses that make the systems brittle and exploitable in specific ways—and vulnerable to failure modalities not seen in traditional software systems. Such failures can rise inadvertently or as the intended results of malicious attacks and manipulation.²⁴ Recent frameworks integrate adversarial attacks²⁵ and unintended faults throughout the lifecycle²⁶ into a single taxonomy that describes both intentional and unintentional failure modes.²⁷

Intentional failures are the result of malicious actors explicitly attacking some aspect of AI system behavior. Taxonomies (e.g., from NIST) on malicious attacks explain the rapidly developing Adversarial Machine Learning (AML) landscape. Attacks span ML training and testing, and each has associated defenses.²⁸ Categories of intentional failures introduced by adversaries include *training data poisoning* attacks (contaminating training data), *model inversion* (recovering training data used in the model through careful queries), and *ML supply chain attacks* (compromising the ML model as it is being downloaded for use).²⁹ National security uses of AI will be the subject of sustained adversarial efforts; AI developed for this community must remain current with a rapidly developing understanding of the nature of vulnerabilities to attacks as these attacks grow in sophistication. Technical and process advances that contribute to reducing vulnerability and to detecting and alerting about attacks must also be monitored routinely.

Unintentional failures can be introduced at any point in the AI development and deployment lifecycle. In addition to faults that can be inadvertently introduced into any software development effort, distinct additional failure modes can be introduced for ML systems.

Examples of unintentional AI failure modes include *reward hacking* (when AI systems learn to achieve a programmed goal in a way that contradicts the programmer's intent) and *distributional shifts* (when a system is tested in one kind of environment but is unable to adapt to changes in other kinds of environments).³⁰ Another area of failure is the inadequate specification of objectives (as described in Section 1 above on *Representing Objectives and Trade-offs*), leading to unexpected and costly behaviors and outcomes.³¹ As AI systems that are separately developed and tested are composed and interact with other AI systems (within one's own services, forces, and agencies, and between U.S. systems and those of allies, adversaries, and potential adversaries), additional unintentional failures can occur.³²

(2) Examples of Current Challenges

To make high-stakes decisions, and often in safety-critical contexts, the DoD and IC must be able to depend on the integrity and security of the data used to train some kinds of ML systems. The challenges of doing so have been echoed by the leadership of the DoD and the IC,³³ including concerns with detecting adversarial attacks such as data poisoning.

(3) Recommendations for Adoption

Critical engineering practices needed to operationalize AI principles (such as “traceable” and “reliable”³⁴) are described in the non-exhaustive list below. These practices span development and fielding of AI systems.

1. *Refine design and development requirements, informed by the concept of operations and risk assessment*, including characterization of failure modes and associated impacts. Conduct systems analysis of operations and identify mission success

metrics and potential functions that can be performed by AI technology. Incorporate early analyses of use cases and scenario development, assess general feasibility and compliance with disallowed outcomes expressed in policy. Critically assess reproducibility (how readily research results can be replicated by a third party) and technical maturity. This includes broad stakeholder engagement and hazard analysis with multidisciplinary experts who ask key questions about potential disparate impacts and document the process undertaken to ensure fairness and the lack of unwanted bias in the ML application.³⁵ The feasibility of meeting these requirements may trigger a review of whether and where it is appropriate to use AI in the system being proposed.

- *Risk assessment.* Trade-offs and risks, including a system's potential societal impact, should be discussed with a diverse, interdisciplinary group. This includes an analysis of the system's potential societal impact and of the impacts of the system's failure modes. Risk-assessment questions should be asked about critical areas relevant to the national security context, including privacy and civil liberties, LOAC, human rights,³⁶ system security, and the risks of a new technology being leaked, stolen, or weaponized.³⁷

2. *Produce documentation of the AI lifecycle.* Whether building and fielding an AI system or "infusing AI" into a preexisting system, require documentation in certain areas.³⁸ These include the data used in ML technologies and the origin of the data³⁹; algorithm(s) used to build models, model characteristics, and intended uses of the AI capabilities; connections between and dependencies within systems, and associated potential complications; the selected testing methodologies, performance indicators, and results for models used in the AI component; and required maintenance (including re-testing requirements) and technical refresh (including for when a system is used in a different scenario/setting or if the AI system is capable of online learning or adaptation).

3. *Leverage infrastructure to support traceability, including auditability and forensics.* Invest resources and build capabilities that support the traceability of AI systems. Traceability captures key information about the system's development and deployment process for relevant personnel to adequately understand the technology.⁴⁰ Audits should support analyses of specific actions and characterizations of longer-term performance and assure that performance on tests of the system and on real-world workloads meet requirements.

4. *For security and robustness, address intentional and unintentional failures.*

- *Adversarial attacks and use of robust ML methods.* Expand notions of adversarial attacks to include various ML attacks⁴¹ (as described above) and seek latest technologies that demonstrate the ability to detect and notify operators of attacks and also tolerate attacks (i.e., to enable systems to withstand or to degrade gracefully when targeted by a deliberate attack).⁴²

- *Follow and incorporate advances in intentional and unintentional ML failures.* Given the rapid evolution of the field of study of intentional and unintentional ML failures, national security organizations must follow and adapt to the latest knowledge about failures and proven practices for system monitoring, failure detection, engineering, and protections during operation. Related efforts and R&D focus on developing and deploying robust AI methods.⁴³
- *Adopt a DevSecOps lifecycle for AI systems focused on potential failure modes.* This includes developing and regularly refining threat models to capture and characterize various attacks, establish a matrixed focus for developing and refining threat models, and ensuring DevSecOps addresses ML development, fielding, and when ML systems are under attack.⁴⁴
- *Limit consequences of system failure through system architecture.* Build an overall system architecture that monitors component performance and handles errors when anomalies are detected; build AI components to be self-protecting and self-checking; and include aggressive stress testing under conditions of intended use.

5. *Conduct red teaming* for both intentional and unintentional failure modalities. Bring together multiple perspectives to rigorously challenge AI systems, exploring the risks, limitations, and vulnerabilities in the context in which they'll be deployed (i.e., red teaming).

- To mitigate intentional failure modes, assume an offensive posture and use methods to make systems more resistant to adversarial attacks, work with adversarial testing tools, and deploy teams dedicated to trying to break systems and make them violate rules for appropriate behavior.⁴⁵
- To mitigate unintentional failure modes, test ML systems per a thorough list of realistic conditions they are expected to operate in. When selecting third-party components, consider the impact that a security vulnerability in them could have on the security of the larger system into which they are integrated. Have an accurate inventory of third-party components and a plan to respond when new vulnerabilities are discovered.
- Organizations should consider establishing broader enterprise-wide communities of AI red teaming capabilities that could be applied to multiple AI developments (e.g., at a DoD service or IC element level, or higher).

(4) Recommendations for Future Action

- *Documentation strategy.* As noted in our First Quarter Recommendations, a common documentation strategy is needed to ensure sufficient documentation by all national security departments and agencies.⁴⁶ In the meantime, agencies should pilot documentation approaches across the AI lifecycle to help inform such a strategy.
- *Standards.* To improve traceability, future work is needed by standard-setting bodies, alongside national security departments/agencies and the broader AI community, to develop audit trail requirements per mission needs for high-stakes AI systems including safety-critical applications (e.g., weapon system controls).
- *Future R&D.* R&D is needed to advance capabilities for cultivating more robust methods that can overcome adverse conditions; to advance approaches that enable assessment of types and levels of vulnerability and immunity; and to tolerate attacks. R&D is also needed to advance capabilities to support risk assessment, including standards, methods, and metrics for evaluating degrees of auditability, traceability, interpretability, explainability, and reliability. For interpretability in particular, R&D is also needed to improve our understanding of the efficacy of interpretability tools and possible interfaces.

III. System Performance

(1) Overview

Fielding AI systems in a responsible manner includes establishing confidence that the technology will perform as intended. An AI system's performance must be assessed,⁴⁷ including assessing its capabilities and blind spots with data representative of real-world scenarios or with simulations of realistic contexts,⁴⁸ and its reliability, robustness (i.e., resilience in real-world settings, including withstanding adversarial attacks on AI components), and security during development and deployment.⁴⁹ System performance must also measure compliance with requirements derived from values such as fairness.

Testing protocols and requirements are essential for measuring and reporting on system performance. (Here, "testing" broadly refers to what the DoD calls "Test and Evaluation, Verification and Validation" [TEVV]. This testing includes both what DoD refers to as Developmental Test and Evaluation and Operational Test and Evaluation.) AI systems present new challenges to established testing protocols and requirements as they increase in complexity, particularly for operational testing. However, existing methods like high-fidelity performance traces and means for sensing shifts (e.g., changes in the statistical distribution of data in operation versus model training) allow for the continuous monitoring of an AI system's performance.

When evaluating system performance, it is especially important to take into account holistic, end-to-end system behavior—the consequence of the interactions and relationships among system elements rather than the independent behavior of individual elements. While system engineering and national security communities have focused on system of systems engineering for years, specific attention must be paid to undesired interactions and emergent performance in AI systems. Multiple relatively independent AI systems can be viewed as distinct agents interacting in the environment of the system of systems, and some of these agents will be humans in and on the loop. Industry has encountered and documented problems in building “systems of systems” out of multiple AI systems.⁵⁰ A related problem is encountered when the performance of one model in a pipeline changes, degrading the overall pipeline behavior.⁵¹ As America’s AI-intensive systems may increasingly be composed and/or interoperable with allied AI-intensive systems, these become important topics for coordination with allies.

(2) Examples of Current Challenges

Unexpected interactions and errors commonly occur in integrated simulations and exercises, illustrating the challenges of predicting and managing behaviors of systems composed of multiple components. Intermittent failures can transpire after composing different systems; these failures are not necessarily the result of any one component having errors, but rather are due to the interactions of the composed systems.⁵²

(3) Recommendations for Adoption

Critical practices to ensure optimal system performance are described in the following non-exhaustive list:

A. Model training and model testing procedures should cover key aspects of performance and appropriate performance metrics.

1. Use regularly updated standards for testing and reporting of system performance. Standards for metrics and reporting are needed to adequately:
 - a. Achieve consistency across testing and test reporting for critical areas.
 - b. Test for blindspots.⁵³
 - c. Test for fairness. When testing for fairness, conduct sustained fairness assessments throughout development and deployment and document deliberations made on the appropriate fairness metrics to use. Agencies should conduct outcome and impact analysis to detect when subtle assumptions in the system show up as unexpected and undesired outcomes in the operational environment.⁵⁴
 - d. Articulate system performance. Clearly document system performance and communicate to the end user the meaning/significance of such performance metrics.

2. *Consider and document the representativeness of the data and model for the specific context at hand.* When using classification and prediction technologies, explicitly consider and document challenges with representativeness of data used in analyses and the fairness/accuracy of inferences and recommendations made with systems leveraging that data when applied in different populations/contexts.

3. *Evaluate an AI system's performance relative to current benchmarks* where possible. Such benchmarks should assist in determining if a proposed AI system's performance meets or exceeds current best performance.

4. *Evaluate aggregate performance of human-machine teams.* Consider that the current benchmark might be the current best performance of a human operator or the composed performance of the human-machine team. Where humans and machines interact, it is important to measure the aggregate performance of the team rather than the AI system alone.⁵⁵

5. *Provide sustained attention to reliability and robustness.* Employ tools and techniques to carefully bound assumptions of robustness of the AI component in the larger system architecture. Provide sustained attention to characterizing the actual performance (for normal and boundary conditions) throughout development and deployment.⁵⁶ For systems of particularly high potential consequences of failure, considerable architecture and design work will have been put into making the overall system fail-safe.

6. *For systems of systems, test machine-machine/multi-agent interaction.* Individual AI systems will be combined in various ways in an enterprise to accomplish broader missions beyond the scope of any single system, which can introduce its own problems.⁵⁷ As a priority during testing, challenge (or "stress test") interfaces and usage patterns with boundary conditions and assumptions about the operational environment and use.

B. Maintenance and deployment

Given the dynamic nature of AI systems, best practices for maintenance are also critically important. Recommended practices include:

1. *Specify maintenance requirements* for datasets as well as for systems, given that their performance can degrade over time.⁵⁸

2. *Continuously monitor and evaluate AI system performance,* including the use of high-fidelity traces to determine continuously if a system is going outside of acceptable parameters.⁵⁹

3. *Conduct iterative model testing and validation.* Training and testing that provide characteristics on capabilities might not transfer or generalize to specific settings of usage; thus, testing and validation may need to be done recurrently, and at strategic intervention points, but especially for new deployments and classes of tasks.⁶⁰

4. *Monitor and mitigate emergent behavior.* There will be instances when systems are composed in ways not anticipated by the developers, thus requiring monitoring the actual performance of the composed system and its components.

(4) *Recommendations for Future Action*

- *Future R&D.* R&D is needed to advance capabilities for TEVV of AI systems to better understand how to conduct persistent and iterative TEVV and build checks and balances into an AI system. Improved methods are needed to explore, predict, and control individual AI system behavior so that when AI systems are composed into systems of systems, their interaction does not lead to unexpected negative outcomes.
- *Metrics.* Progress on a common understanding of TEVV concepts and requirements is critical for progress in widely used metrics for performance. Significant work is needed to establish what appropriate metrics should be used to assess system performance across attributes for responsible AI according to applications/context profiles. (Such attributes, for example, include fairness, interpretability, reliability, and robustness.) Future work is needed to develop: (1) definitions, taxonomy, and metrics needed to enable agencies to better assess AI performance and vulnerabilities; and (2) metrics and benchmarks to assess reliability and intelligibility of produced model explanations. In the near term, guidance is needed on: (1) standards for testing intentional and unintentional failure modes; (2) exemplar data sets for benchmarking and evaluation, including robustness testing and red teaming; and (3) defining characteristics of AI data quality and training environment fidelity (to support adequate performance and governance).⁶¹
- *International collaboration and cooperation.* Collaboration is needed to align on how to test and verify AI system reliability and performance, including along shared values (such as fairness and privacy). Such collaboration will be critical among allies and partners for interoperability and trust. Additionally, these efforts could potentially include dialogues between the U.S. and strategic competitors on establishing common standards of AI safety and reliability testing to reduce the chances of inadvertent escalation.

IV. Human-AI Interaction & Teaming

(1) *Overview*

Responsible AI development and fielding requires striking the right balance of leveraging human and AI reasoning, recommendation, and decision-making processes. Ultimately,

all AI systems will have some degree of human-AI interaction as they all will be developed to support humans. And some systems will serve as more than just support tools and will adopt roles of teammates that actively collaborate with humans.

(2) Examples of Current Challenges

There is an opportunity to develop AI systems to complement and augment human understanding, decision-making, and capabilities. Decisions about developing and fielding AI systems for specific domains or scenarios should consider the relative strengths of AI capabilities and human intellect across the expected range of tasks, considering AI system maturity or capability and how people and machines might coordinate.

Designs and methods for human-AI interaction can be employed to enhance human-AI teaming.⁶² Methods in support of effective human-AI interaction can help AI systems understand when and how to engage humans for assistance, when AI systems should take initiative to assist human operators, and, more generally, how to support the creation of effective human-AI teams. In engaging with end users, it may be important for AI systems to infer and share with end users well-calibrated levels of confidence about their inferences, to provide human operators with an ability to weigh the importance of machine output or pause to consider details behind a recommendation more carefully. Methods, representations, and machinery can be employed to provide insight about AI inferences, including the use of interpretable machine learning.⁶³

Research directions include developing and fielding machinery aimed at reasoning about human strengths and weaknesses, such as recognizing and responding to the potential for costly human biases of judgment and decision-making in specific settings.⁶⁴ Other work centers on mechanisms to consider the ideal mix of initiatives, including when and how to rely on human expertise versus on AI inferences.⁶⁵ As part of effective teaming, AI systems can be endowed with the ability to detect the focus of attention, workload, and sensitivity to interruption of human operators and consider these inferences in decisions about when and how to engage with operators.⁶⁶ Directions of effort include developing mechanisms for identifying the most relevant information or inferences to provide end users with different skill levels in different settings.⁶⁷ Consideration must be given to the prospect of introducing bias, including potential biases that may arise because of the configuration and sequencing of rendered data. For example, IC research⁶⁸ shows that confirmation bias can be triggered by the order in which information is displayed, and this order can consequently impact or sway intel analyst decisions. Careful design and study can help to identify and mitigate such bias.

(3) Recommendations for Adoption

Critical practices to ensure optimal human-AI interaction are described in the non-exhaustive list below. These recommended practices span the entire AI lifecycle.

A. Identification of functions of humans in design, engineering, and fielding of AI.

1. Given AI and human capabilities and complementarities, as well as requirements for accountability and human judgment, define the tasks of humans and the goals and mission of the human-machine team across the AI lifecycle. This entails noting needs for feedback loops, including opportunities for oversight.

2. Define functions and responsibilities of humans during system operation and assign them to specific individuals. Functions and responsibilities will vary for each domain and project and should be periodically revisited.

B. Explicit support of human-AI interaction and collaboration.

1. *Extend human-AI design methodologies and guidelines.* AI systems designs should take into account the defined tasks of humans in human-AI collaborations in different scenarios; ensure that the mix of human-machine actions in the aggregate is consistent with the intended behavior and accounts for the ways that human and machine behavior can co-evolve⁶⁹; and also avoid automation bias (that places unjustified confidence in the results of the computation) and unjustified reliance on humans in the loop as fail-safe mechanisms. Practices should allow for auditing of the human-AI pair and designs should be transparent to allow for an understanding of how the AI is working day-to-day, supported by an audit trail if things go wrong. Based on context and mission need, designs should ensure usability of AI systems by AI experts, domain experts, and novices, as appropriate.

2. *Employ algorithms and functions in support of interpretability and explanation.* Algorithms and functions that provide individuals with task-relevant knowledge and understanding should take into account that key factors in an AI system's inferences and actions can be understood differently by various audiences (e.g., real-time operators, engineers and data scientists, and oversight officials). Interpretability and explainability exists in degrees. In this regard, interpretability intersects with traceability, audit, and documentation practices.

3. *Design systems to provide cues to human operator(s) about the level of confidence the system has in its results or behaviors.* AI system designs should appropriately convey uncertainty and error bounding. For instance, a user interface should convey system self-assessment of confidence alerts when the operational environment is significantly different from the environment the system was trained for and indicate internal inconsistencies that call for caution.

4. *Refine policies for machine-human initiative and handoff.* Policies, and aspects of human-computer interaction, system interface, and operational design, should define when and how information or tasks should be passed from a machine to a human operator and vice versa.

5. *Leverage traceability to assist with system development and understanding.* Traceability processes must capture details about human-AI interaction to retroactively understand where challenges occurred, and why, in order to improve systems and their use for redress. Infrastructure and instrumentation⁷⁰ can also help assess humans, systems, and environments to gauge the impact of AI at all levels of system maturity and to measure the effectiveness and performance for hybrid human-AI systems in a mission context.

6. *Conduct training.* Train and educate individuals responsible for AI development and fielding, including human operators, decision-makers, and procurement officers.⁷¹

(4) Recommendations for Future Action

- *Future R&D.* R&D is needed to advance capabilities of AI technologies to perceive and understand the meaning of human communication, including spoken speech, written text, and gestures. This research should account for varying languages and cultures, with special attention to diversity given that AI often performs worse in cases impacting gender and racial minorities. It is also needed to improve human-machine teaming, including disciplines and technologies centered on decision sciences, control theory, psychology, economics (human aspects and incentives), and human factors engineering. R&D for human-machine teaming should also focus on helping systems understand human blind spots and biases and optimizing factors such as human attention, human workload, ideal mixing of human and machine initiative, and passing control between the human and machine. R&D also is needed to optimize the ability of humans and AI to work together to undertake complex, evolving tasks in a variety of environments, as well as for diverse groupings of machines to cooperate with each other, with broader systems, and with human counterparts to achieve shared objectives.

- *Training.* Ongoing work is needed to train the workforce that will interact with, collaborate with, and be supported by AI systems. In its First Quarter Recommendations, the Commission provided recommendations for such training. Operators should receive training on the specifics of the system and application, the fundamentals of AI and data science, and refresher trainings (e.g., when systems are deployed in new settings and unfamiliar scenarios, and when predictive models are revised with new data, as performance may shift with updates and introduce behaviors unfamiliar to operators).

V. Accountability and Governance

(1) Overview

National security departments and agencies must specify who will be held accountable for both specific system outcomes and general system maintenance and auditing, in what way, and for what purpose. Government must address the difficulties in preserving human accountability, including for end users, developers, testers, and the organizations employing AI systems. End users and those affected by the actions of an AI system should have the opportunity to appeal an AI system's determinations. Accountability and appellate processes must exist for AI decisions, inferences, recommendations, and actions.

(2) Examples of Current Challenges

If a contentious outcome occurs, overseeing entities need the technological capacity to understand what in the AI system caused this. For example, if a soldier uses an AI-enabled weapon and the result violates international law of war standards, an investigating body or military tribunal should be able to re-create what happened through audit trails and other documentation. Without policies requiring such technology and the enforcement of those policies, proper accountability would be elusive, if not impossible. Moreover, auditing trails and documentation will prove critical as courts begin to grapple with whether AI system determinations reach the requisite standards to be admitted as evidence. Building the traceability infrastructure to permit auditing (as described in *Engineering Practices*) will increase the costs of building AI systems and take significant work—a necessary investment given our commitment to accountability, discoverability, and legal compliance.

(3) Recommendations for Adoption

Critical accountability and governance practices are identified in the non-exhaustive list below.

1. *Appoint full-time responsible AI leads* to join senior leadership. Every department and agency critical to national security and each branch of the armed services, at a minimum, should have a dedicated, full-time responsible AI lead who is part of the senior leadership team. Such leads should oversee the implementation of the Key Considerations recommended practices alongside the department or agency's respective AI principles.

2. *Identify responsible actors.* Determine and document the people accountable for a specific AI system or any given part of the system and the processes involved. This includes identifying who is responsible for the development or procurement; operation (including the system's inferences, recommendations, and actions during usage), and maintenance of an AI system, as well as the authorization of a system and enforcement of policies for use. Determine and document the mechanism/structure for holding such actors accountable and to whom it should be disclosed for proper oversight.

3. *Require technology to strengthen accountability processes and goals.* Document the chains of custody and command involved in developing and fielding AI systems to know who was responsible at which point in time. Improving traceability and auditability capabilities will allow agencies to better track a system's performance and outcomes.⁷² Policy should establish requirements about information that should be captured about the development process and about system performance and behavior in operation.

4. *Adopt policies to strengthen accountability and governance.* Identify or, if lacking, establish policies that allow individuals to raise concerns about irresponsible AI development/fielding (e.g., via an ombudsman). This requires ensuring a governance structure is in place to address grievances and harms if systems fail, which supports feedback loops and oversight to ensure that systems operate as they should.

Agencies should institute specific oversight and enforcement practices, including auditing and reporting requirements; a mechanism that would allow thorough review of the most sensitive/high-risk AI systems to ensure auditability and compliance with responsible use and fielding requirements; an appealable process for those found at fault for developing or using AI irresponsibly; and grievance processes for those affected by the actions of AI systems. Agencies should leverage best practices from academia and industry for conducting internal audits and assessments,⁷³ while also acknowledging the benefits offered by external audits.⁷⁴

5. *Support external oversight.* Remain responsive and facilitate oversight through documentation processes and other policy decisions.⁷⁵ For instance, supporting traceability and specifically documentation to audit trails will allow for external oversight.⁷⁶ Self-assessment alone might prove to be inadequate in all scenarios.⁷⁷ Congress can provide a key oversight function throughout the AI lifecycle, asking critical questions of agency leaders and those responsible for AI systems.

(4) Recommendations for Future Action

Currently no external oversight mechanism exists specific to AI in national security. Notwithstanding the important work of Inspectors General in conducting internal oversight, open questions remain as to how to complement current practices and structures.

Appendix C - Endnotes

¹ Examples of efforts to establish ethics guidelines are found within the U.S. government, industry, and internationally. See, e.g., *Draft Memorandum for the Heads of Executive Departments and Agencies: Guidance for Regulation of Artificial Intelligence Applications*, Office of Management and Budget (Jan. 1, 2019), <https://www.whitehouse.gov/wp-content/uploads/2020/01/Draft-OMB-Memo-on-Regulation-of-AI-1-7-19.pdf>; Jessica Fjeld & Adam Nagy, *Principled Artificial Intelligence: Mapping Consensus in Ethical and Rights-Based Approaches to Principles for AI*, Berkman Klein Center (Jan. 15, 2020), <https://cyber.harvard.edu/publication/2020/principled-ai>; *OECD Principles on AI*, OECD (last visited June 17, 2020), <https://www.oecd.org/going-digital/ai/principles/>; *Ethics Guidelines for Trustworthy AI*, European Commission at 26-31 (April 8, 2019), <https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai>; *Assessment List for Trustworthy Artificial Intelligence (ALTAI) for Self-assessment*, European Commission (July 17, 2020), <https://ec.europa.eu/digital-single-market/en/news/assessment-list-trustworthy-artificial-intelligence-altai-self-assessment>.

² C. Todd Lopez, *DOD Adopts 5 Principles of Artificial Intelligence Ethics*, U.S. Department of Defense (Feb. 5, 2020), <https://www.defense.gov/Explore/News/Article/Article/2094085/dod-adopts-5principles-of-artificial-intelligence-ethics/> [hereinafter Lopez, DoD Adopts 5 Principles].

³ See Ben Huebner, *Presentation: AI Principles*, Intelligence and National Security Alliance 2020 Spring Symposium: Building an AI-Powered IC (March 4, 2020), <https://www.insaonline.org/2020-spring-symposium-building-an-ai-powered-ic-event-recap/>.

⁴ See, e.g., U.S. Const. amendments I, IV, V, and XIV; Americans with Disabilities Act of 1990, 42 U.S.C. § 12101 et seq.; Title VII of the Consumer Credit Protection Act, 15 U.S.C. §§ 1691-1691f; Title VII of the Civil Rights Act of 1964, 42 U.S.C. § 2000e et seq.

⁵ International Covenant on Civil and Political Rights, UN General Assembly, United Nations, Treaty Series, vol. 999, at 171 (Dec. 16, 1966), <https://www.refworld.org/docid/3ae6b3aa0.html>. As noted in the Commission's *Interim Report*, America and its like-minded partners share a commitment to democracy, human dignity, and human rights. *Interim Report*, NSCAI (Nov. 2019), <https://www.nscai.gov/previous-reports/>. Many, but not all nations, share commitments to these values. Even when values are shared, however, they can be culturally relative, for instance, across nations, owing to interpretative nuances.

⁶ See, e.g., Daniel Coats, *Intelligence Community Directive 107*, Office of the Director of National Intelligence (Feb. 28, 2018), <https://fas.org/irp/dni/icd/icd-107.pdf> (on protecting civil liberties and privacy); *IC Framework for Protecting Civil Liberties and Privacy and Enhancing Transparency Section 702*, Intel.gov (Jan. 2020), https://www.intelligence.gov/index.php/ic-on-the-record/guide-to-posted-documents#SECTION_702-OVERVIEW (on privacy and civil liberties implication assessments and oversight); *Principles of Professional Ethics for the Intelligence Community*, Office of the Director of National Intelligence (last accessed June 17, 2020), <https://www.dni.gov/index.php/who-we-are/organizations/clpt/clpt-related-menus/clpt-related-links/ic-principles-of-professional-ethics> (on diversity and inclusion).

⁷ See, e.g., *Privacy Office*, U.S. Department of Homeland Security (last accessed June 3, 2020), <https://www.dhs.gov/privacy-office#>; *CRCL Compliance Branch*, U.S. Department of Homeland Security (last accessed May 15, 2020), <https://www.dhs.gov/compliance-branch>.

⁸ See Samuel Jenkins & Alexander Joel, *Balancing Privacy and Security: The Role of Privacy and Civil Liberties in the Information Sharing Environment*, IAPP Conference 2010 (2010), <https://dpcl.dod.defense.gov/Portals/49/Documents/Civil/IAPP.pdf>.

⁹ See *Projects*, U.S. Privacy and Civil Liberties Oversight Board (last visited June 17, 2020), <https://www.pclob.gov/Projects>.

¹⁰ See *Department of Defense Law of War Manual*, U.S. Department of Defense (Dec. 2016), <https://dod.defense.gov/Portals/1/Documents/pubs/DoD%20Law%20of%20War%20Manual%20-%20June%202015%20Updated%20Dec%202016.pdf?ver=2016-12-13-172036-190> [hereinafter DoD Law of War Manual]; see also *AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense: Supporting Document*, DoD Defense Innovation Board (Oct. 31, 2019), https://media.defense.gov/2019/Oct/31/2002204459/-1/-1/0/DIB_AI_Principles_supporting_document.pdf ("More than 10,000 military and civilian lawyers within DoD advise on legal compliance with regard to the entire range of DoD activities, including the Law of War. Military lawyers train DoD personnel on Law of War requirements, for example, by providing additional Law of War instruction prior to a deployment of forces abroad. Lawyers for a Component DoD organization advise on the

issuance of plans, policies, regulations, and procedures to ensure consistency with Law of War requirements. Lawyers review the acquisition or procurement of weapons. Lawyers help administer programs to report alleged violations of the Law of War through the chain of command and also advise on investigations into alleged incidents and on accountability actions, such as commanders' decisions to take action under the Uniform Code of Military Justice. Lawyers also advise commanders on Law of War issues during military operations.”).

¹¹ Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment, United Nations General Assembly (Dec. 10, 1984), <https://www.ohchr.org/en/professionalinterest/pages/cat.aspx>.

¹² See DoD Law of War Manual at 26 (“Rules of Engagement reflect legal, policy, and operational considerations, and are consistent with the international law obligations of the United States, including the law of war.”).

¹³ See *Department of Defense Directive 3000.09 on Autonomy in Weapon Systems*, U.S. Department of Defense (Nov. 21, 2012), <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodd/300009p.pdf> (“Autonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force.”).

¹⁴ See, e.g., *Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System*, Partnership on AI, <https://www.partnershiponai.org/report-on-machine-learning-in-risk-assessment-tools-in-the-u-s-criminal-justice-system/>; Jeffrey Dastin, *Amazon Scraps Secret AI Recruiting Tool that Showed Bias Against Women*, Reuters (Oct. 10, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazonscraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G> [*hereinafter* Dastin, Amazon Scraps Secret AI Recruiting Tool]; Andi Peng et al., *What You See Is What You Get? The Impact of Representation Criteria on Human Bias in Hiring*, Proceedings of the 7th AAAI Conference on Human Computation and Crowdsourcing (Oct. 2019), <https://arxiv.org/pdf/1909.03567.pdf>; Patrick Grother, et al., *Face Recognition Vendor Test (FRVT) Part Three: Demographic Effects*, National Institute of Standards and Technology (Dec. 2019), <https://doi.org/10.6028/NIST.IR.8280>.

¹⁵ PNDC provides predictive analytics to improve military readiness; enable earlier identification of service members with potential unfitting, disabling, or career-ending conditions; and offer opportunities for early medical intervention or referral into disability processing. To do so, PNDC provides recommendations at multiple points in the journey of the non-deployable service member through the Military Health System to make “better decisions” that improve medical outcomes and delivery of health services. This is very similar to the OPTUM decision support system that recommended which patients should get additional intervention to reduce costs. Analysis showed millions of U.S. patients were processed by the system, with substantial disparate impact on Black patients compared to white patients. Shaping development from the start to reflect bias issues (which can be subtle) would have produced a more equitable system and avoided scrutiny and suspension of system use when findings were disclosed. Heidi Ledford, *Millions of Black People Affected by Racial Bias in Health Care Algorithms*, Nature (Oct. 26, 2019), <https://www.nature.com/articles/d41586-019-03228-6>.

¹⁶ See e.g., Dastin, Amazon Scraps Secret AI Recruiting Tool.

¹⁷ This combined approach of stable policy-level disallowed outcomes and system-specific disallowed outcomes is consistent with DoD practices for system safety, for example. See *Department of Defense Standard Practice: System Safety*, U.S. Department of Defense (May 11, 2012), <https://www.dau.edu/cop/armysoh/DAU%20Sponsored%20Documents/MIL-STD-882E.pdf>. Depending on the context, mitigating harm per values and disallowed outcomes might entail the use of fail-safe technologies. See Eric Horvitz, *Reflections on Safety and Artificial Intelligence, Exploratory Technical Workshop on Safety and Control for AI* (June 27, 2016), http://erichorvitz.com/OSTP-CMU_AI_Safety_framing_talk.pdf. See also Dorsa Sadigh & Ashish Kapoor, *Safe Control Under Uncertainty with Probabilistic Signal Temporal Logic*, Proceedings of Robotics: Science and Systems XII (2016), <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/11/RSS2016.pdf>.

¹⁸ Mohsen Bayati, et al., *Data-Driven Decisions for Reducing Readmissions for Heart Failure: General Methodology and Case Study*, PLOS One Medicine (Oct. 8, 2014), <https://doi.org/10.1371/journal.pone.0109264>; Eric Horvitz & Adam Seiver, *Time-Critical Action: Representations and Application*, Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (Aug. 1997), <https://arxiv.org/pdf/1302.1548.pdf>.

Appendix C - Endnotes

¹⁹ See Inioluwa Deborah Raji, et al., *Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing*, ACM FAT (Jan. 3, 2020), <https://arxiv.org/abs/2001.00973> [hereinafter Raji, Closing the AI Accountability Gap].

²⁰ See Lopez, DoD Adopts 5 Principles.

²¹ *Model Interpretability in Azure Machine Learning*, Microsoft (Nov. 16, 2020), <https://docs.microsoft.com/en-us/azure/machine-learning/how-to-machine-learning-interpretability>.

²² Lopez, DoD Adopts 5 Principles.

²³ Jessica Cussins Newman, *Decision Points in AI Governance: Three Case Studies Explore Efforts to Operationalize AI Principles* (May 5, 2020), Berkeley Center for Long-Term Cybersecurity, <https://cltc.berkeley.edu/ai-decision-points/>; Raji, *Closing the AI Accountability Gap*; Miles Brundage, et al., *Toward Trustworthy AI Development: Mechanisms for Supporting Verifiable Claims* (April 20, 2020), <https://arxiv.org/abs/2004.07213> [hereinafter Brundage, *Toward Trustworthy AI Development*]; Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, Microsoft (March 2019), https://www.microsoft.com/en-us/research/uploads/prod/2019/03/amershi-icse-2019-Software_Engineering_for_Machine_Learning.pdf.

²⁴ Dario Amodei, et al., *Concrete Problems in AI Safety*, arXiv (July 25, 2016), <https://arxiv.org/abs/1606.06565>.

²⁵ Guofu Li, et al., *Security Matters: A Survey on Adversarial Machine Learning*, arXiv (Oct. 23, 2018), <https://arxiv.org/abs/1810.07339>; Elham Tabassi et al., *NISTIR 8269: A Taxonomy and Terminology of Adversarial Machine Learning (Draft)*, National Institute of Standards and Technology (Oct. 2019), <https://csrc.nist.gov/publications/detail/nistir/8269/draft>.

²⁶ José Faria, *Non-Determinism and Failure Modes in Machine Learning*, 2017 IEEE 28th International Symposium on Software Reliability Engineering Workshops (Oct. 2017), <https://ieeexplore.ieee.org/document/8109300>.

²⁷ Ram Shankar Siva Kumar, et al. *Failure Modes in Machine Learning* (Nov. 11, 2019), <https://docs.microsoft.com/en-us/security/engineering/failure-modes-in-machine-learning> [hereinafter Kumar, *Failure Modes in Machine Learning*].

²⁸ See Elham Tabassi et al., *NISTIR 8269: A Taxonomy and Terminology of Adversarial Machine Learning (Draft)*, National Institute of Standards and Technology (Oct. 2019), <https://csrc.nist.gov/publications/detail/nistir/8269/draft>. See also Kumar, *Failure Modes in Machine Learning*.

²⁹ For 11 categories of attack, and associated overviews, see the Intentionally-Motivated Failures Summary in Kumar, *Failure Modes in Machine Learning*.

³⁰ For more on reward hacking, see Jack Clark, et al., *Faulty Reward Functions in the Wild* (Dec. 21, 2016), <https://openai.com/blog/faulty-reward-functions/>. For more on distributional shifts, see Colin Smith, et al., *Hazard Contribution Modes of Machine Learning Components*, AAAI-20 Workshop on Artificial Intelligence Safety (SafeAI 2020) (Feb. 7, 2020), <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20200001851.pdf> (Unexpected performance represents emergent runtime output, behavior, or effects at the system level, e.g., through unanticipated feature interaction ... that was also not previously observed during model validation.).

³¹ Thomas Dietterich & Eric Horvitz, *Rise of Concerns About AI: Reflections and Directions*, Communications of the ACM at 38-40 (Oct. 2015), http://erichorvitz.com/CACM_Oct_2015-VP.pdf. See also Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, Science (Oct. 25, 2019), <https://science.sciencemag.org/content/366/6464/447>.

³² Kumar, *Failure Modes in Machine Learning*.

³³ For concerns about generative adversarial networks (GANS) voiced by Gen. Shanahan, JAIC, see Don Rassler, *A View from the CT Foxhole: Lieutenant General John N.T. "Jack" Shanahan, Director, Joint Artificial Intelligence Center, Department of Defense*, Combating Terrorism Center at West Point (Dec. 2019), <https://ctc.usma.edu/view-ct-foxhole-lieutenant-general-john-n-t-jack-shanahan-director-joint-artificial-intelligence-center-department-defense/>. Concerns about GANS, information authenticity, and reliable and understandable systems were voiced by Dean Souleles, IC. See *Afternoon Keynote*, Intelligence and National Security Alliance 2020 Spring Symposium: Building an AI-Powered IC (March 4, 2020), <https://www.insonline.org/2020-spring-symposium-building-an-ai-powered-ic-event-recap/>.

³⁴ See Lopez, DOD Adopts 5 Principles.

³⁵ There is no single definition of fairness. System developers and organizations fielding applications must work with stakeholders to define fairness and provide transparency via disclosure of assumed definitions of fairness. Definitions or assumptions about fairness and metrics for identifying fair inferences and allocations should be explicitly documented. This should be accompanied by a discussion of alternate definitions and rationales for the current choice. These elements should be documented internally as ML components and larger systems are developed. This is especially important as establishing alignment on the metrics to use for assessing fairness encounters an added challenge when different cultural and policy norms are involved when collaborating on development and use with allies.

³⁶ For more on the importance of human rights impact assessments of AI systems, see *Report of the Special Rapporteur to the General Assembly on AI and Its Impact on Freedom of Opinion and Expression*, UN Human Rights Office of the High Commissioner (Aug. 29, 2018), <https://www.ohchr.org/EN/Issues/FreedomOpinion/Pages/ReportGA73.aspx>. For an example of a human rights risk assessment for AI in categories such as nondiscrimination and equality, political participation, privacy, and freedom of expression, see Mark Latonero, *Governing Artificial Intelligence: Upholding Human Rights & Dignity*, Data Society (Oct. 2018), https://datasociety.net/wp-content/uploads/2018/10/DataSociety_Governing_Artificial_Intelligence_Upholding_Human_Rights.pdf.

³⁷ For exemplary risk assessment questions that IARPA has used, see Richard Danzig, *Technology Roulette: Managing Loss of Control as Many Militaries Pursue Technological Superiority*, Center for a New American Security at 22 (June 28, 2018), <https://s3.amazonaws.com/files.cnas.org/documents/CNASReport-Technology-Roulette-DoSproof2v2.pdf?mtime=20180628072101>.

³⁸ Documentation recommendations build off of a legacy of robust documentation requirements. See *Department of Defense Standard Practice: Documentation of Verification, Validation, and Accreditation (VV&A) For Models and Simulations*, Department of Defense (Jan. 28, 2008), <https://acqnotes.com/Attachments/MIL-STD-3022%20Documentation%20of%20VV&A%20for%20Modeling%20%20Simulation%2028%20Jan%2008.pdf>.

³⁹ For an industry example, see Timnit Gebru, et al., *Datasheets for Datasets*, Microsoft (March 2018), <https://www.microsoft.com/en-us/research/publication/datasheets-for-datasets/>. For more on data, model, and system documentation, see *Annotation and Benchmarking on Understanding and Transparency of Machine Learning Lifecycles (ABOUT ML)*, an evolving body of work from the Partnership on AI about documentation practices at <https://www.partnershiponai.org/about-ml/>. Documenting caveats of re-use for both data sets and models is critical to avoid "off-label" use harms, as one senior official notes. David Thornton, *Intelligence Community Laying Foundation for AI Data Analysis*, Federal News Network (Nov. 1, 2019), <https://federalnewsnetwork.com/allnews/2019/11/intelligence-community-laying-the-foundation-for-ai-data-analysis/>.

⁴⁰ Jonathan Mace, et al., *Pivot Tracing: Dynamic Causal Monitoring for Distributed Systems*, Communications of the ACM, Vol. 63 No. 3, at 94-102 (March 2020), <https://m-cacm.acm.org/magazines/2020/3/243034-pivot-tracing/fulltext> [hereinafter Mace, Pivot Tracing].

⁴¹ Aleksander Madry, et al., *Towards Deep Learning Models Resistant to Adversarial Attacks*, MIT (Sept. 4, 2019), <https://arxiv.org/abs/1706.06083> [hereinafter Madry, Towards Deep Learning Models Resistant to Adversarial Attacks].

⁴² See e.g., *id.*; Thomas Dietterich, *Steps Toward Robust Artificial Intelligence*, Association for the Advancement of Artificial Intelligence (Fall 2017), <https://www.aaai.org/ojs/index.php/aimagazine/article/view/2756/2644>; Eric Horvitz, *Reflections on Safety and Artificial Intelligence* (June 27, 2016), http://erichorvitz.com/OSTP-CMU_AI_Safety_framing_talk.pdf.

Appendix C - Endnotes

⁴³ On adversarial attacks on ML, see Kevin Eykholt, et al., *Robust Physical-World Attacks on Deep Learning Visual Classification*, IEEE Conference on Computer Vision and Pattern Recognition at 1625-1634 (June 18-23, 2018), <https://ieeexplore.ieee.org/document/8578273>. On directions with robustness, see Madry, *Towards Deep Learning Models Resistant to Adversarial Attacks*. For a more exhaustive list of sources see *Key Considerations for Responsible Development & Fielding of Artificial Intelligence: Extended Version*, NSCAI (2021) (on file with the Commission).

⁴⁴ Ram Shankar Siva Kumar, et al., *Adversarial Machine Learning—Industry Perspectives*, 2020 IEEE Symposium on Security and Privacy (SP) Deep Learning and Security Workshop (May 21, 2020), <https://arxiv.org/abs/2002.05646>.

⁴⁵ Dou Goodman, et al., *Advbox: A Toolbox to Generate Adversarial Examples That Fool Neural Networks* (Aug. 26, 2020), <https://arxiv.org/abs/2001.05574>.

⁴⁶ See *First Quarter Recommendations*, NSCAI (March 2020), <https://www.nscai.gov/previous-reports/>. Ongoing efforts to share best practices for documentation among government agencies through GSA's AI Community of Practice further indicate the ongoing need and desire for common guidance.

⁴⁷ Ben Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*, International Journal of Human-Computer Interaction 2020 at 495-504 (March 23, 2020), <https://doi.org/10.1080/10447318.2020.1741118> [*hereinafter* Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*].

⁴⁸ However, test protocols must acknowledge that test sets may not be fully representative of real-world usage.

⁴⁹ Brundage, *Toward Trustworthy AI Development*; Ece Kamar, et al., *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*, Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (June 2012), <https://dl.acm.org/doi/10.5555/2343576.2343643> [*hereinafter* Kamar, *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*].

⁵⁰ One example is “Hidden Feedback Loops,” where systems that learn from external-world behavior may also shape the behavior they are monitoring. See D. Sculley, et al., *Machine Learning: The High Interest Credit Card of Technical Debt*, Google (2014), <https://research.google/pubs/pub43146/>.

⁵¹ Megha Srivastava, et al., *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*, KDD'20 (Aug. 11, 2020), <https://arxiv.org/abs/2008.04572> [*hereinafter* Srivastava, *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*].

⁵² David Sculley, et al., *Hidden Technical Debt in Machine Learning Systems*, Proceedings of the 28th International Conference on Neural Information Processing Systems (Dec. 2015), <https://dl.acm.org/doi/10.5555/2969442.2969519>.

⁵³ Ramya Ramakrishnan, et al., *Blind Spot Detection for Safe Sim-to-Real Transfer*, Journal of Artificial Intelligence Research 67 at 191-234 (Feb. 4, 2020), <https://www.jair.org/index.php/jair/article/view/11436>.

⁵⁴ See Microsoft's AI Fairness checklist as an example of an industry tool to support fairness assessments; Michael A. Madaio, et al., *Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI*, CHI 2020 (April 25-30, 2020), <http://www.jennvw.com/papers/checklists.pdf> [*hereinafter* Madaio, *Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI*].

⁵⁵ Kamar, *Combining Human and Machine Intelligence in Large-Scale Crowdsourcing*.

⁵⁶ See Shneiderman, *Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy*.

⁵⁷ Cynthia Dwork, et al., *Individual Fairness in Pipelines*, arXiv (April 12, 2020), <https://arxiv.org/abs/2004.05167>; Srivastava, *An Empirical Analysis of Backward Compatibility in Machine Learning Systems*.

⁵⁸ *Artificial Intelligence (AI) Playbook for the U.S. Federal Government*, Artificial Intelligence Working Group, ACT-IAC Emerging Technology Community of Interest (Jan. 22, 2020), <https://www.actiac.org/act-iac-white-paper-artificial-intelligence-playbook>.

⁵⁹ Ori Cohen, *Monitor! Stop Being A Blind Data-Scientist*, Towards Data Science (Oct. 8, 2019), <https://towardsdatascience.com/monitor-stop-being-a-blind-data-scientist-ac915286075f>; Mace, Pivot Tracing at 94-102.

⁶⁰ Eric Breck, et al., *The ML Test Score: A Rubric for ML Production Readiness and Technical Debt Reduction*, 2017 IEEE International Conference on Big Data (Dec. 11-14, 2017), <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=8258038&tag=1>.

⁶¹ The 2021 NDAA expansion of the National Institute of Standards & Technology (NIST) mission authorizes the standards body to provide such guidance: “National Institute of Standards and Technology Activities (Title LIII, Sec. 5301)—expands NIST mission to include advancing collaborative frameworks, standards, guidelines for AI, supporting the development of a risk-mitigation framework for AI systems, and supporting the development of technical standards and guidelines to promote trustworthy AI systems.” Pub. L. 116-283, William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, 134 Stat. 3388 (2021).

⁶² Saleema Amershi, et al., *Guidelines for Human-AI Interaction*, CHI '19: Proceedings of the CHI Conference on Human Factors in Computing Systems (May 2019), <https://dl.acm.org/doi/10.1145/3290605.3300233>.

⁶³ Rich Caruana, et al., *Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission*, Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Aug. 10-13, 2015), <https://www.semanticscholar.org/paper/Intelligible-Models-for-HealthCare%3APredicting-Risk-Caruana-Lou/cb030975a3dbcdf52a01cbd1c140711332313e13>.

⁶⁴ Eric Horvitz, *Reflections on Challenges and Promises of Mixed-Initiative Interaction*, AI Magazine (Summer 2007), http://erichorvitz.com/mixed_initiative_reflections.pdf.

⁶⁵ Eric Horvitz, *Principles of Mixed-Initiative User Interfaces*, Proceedings of CHI '99 ACM SIGCHI Conference on Human Factors in Computing Systems (May 1999), <https://dl.acm.org/doi/10.1145/302979.303030>; Kamar, Combining Human and Machine Intelligence in Large-Scale Crowdsourcing.

⁶⁶ Eric Horvitz, et al., *Models of Attention in Computing and Communications: From Principles to Applications*, Communications of the ACM at 52-59 (March 2003), <https://cacm.acm.org/magazines/2003/3/6879-models-of-attention-in-computingand-communication/fulltext>.

⁶⁷ Eric Horvitz & Matthew Barry, *Display of Information for Time-Critical Decision Making*, Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (Aug. 1995), <https://arxiv.org/pdf/1302.4959.pdf>.

⁶⁸ There has been considerable research in the IC on the challenges of confirmation bias for analysts. Some experiments demonstrated a strong effect that the sequence in which information is presented alone can shape analyst interpretations and hypotheses. Brant Cheikes, et al., *Confirmation Bias in Complex Analyses*, MITRE (Oct. 2004), https://www.mitre.org/sites/default/files/pdf/04_0985.pdf. This highlights the care that is required when designing the human-machine teaming when complex, critical, and potentially ambiguous information is presented to analysts and decision-makers.

⁶⁹ Shneiderman, Human-Centered Artificial Intelligence: Reliable, Safe & Trustworthy at 495-504. An example of co-evolution of machine and human behavior is in ML spam filters. As human spammers determine what characteristics are getting email flagged as spam, they change how they generate spam, which requires the spam-detection models to evolve in a constant “arms race.”

⁷⁰ Infrastructure includes tools (hardware and software) in the test environment that support monitoring system performance (such as the timing of exchanges among systems or the ability to generate test data). Instrumentation refers to the presence of monitoring and additional interfaces to provide insight into a specific system under test.

⁷¹ Jamie Berryhill, et al., *Hello, World: Artificial Intelligence and Its Use in the Public Sector*, OECD Working Papers on Public Governance (Nov. 21, 2019), <https://doi.org/10.1787/726fd39d-en>.

Appendix C - Endnotes

⁷² See Raji, Closing the AI Accountability Gap.

⁷³ See *Id.* (“In this paper, we present internal algorithmic audits as a mechanism to check that the engineering processes involved in AI system creation and deployment meet declared ethical expectations and standards, such as organizational AI principles”); see also Madaio, Co-Designing Checklists to Understand Organizational Challenges and Opportunities Around Fairness in AI.

⁷⁴ For more on the benefits of external audits, see Brundage, Toward Trustworthy AI Development. For an agency example, see Aaron Boyd, *CBP Is Upgrading to a New Facial Recognition Algorithm in March*, Nextgov.com (Feb. 7, 2020), <https://www.nextgov.com/emerging-tech/2020/02/cbp-upgrading-new-facialrecognition-algorithm-march/162959/> (highlighting a NIST algorithmic assessment on behalf of U.S. Customs and Border Protection).

⁷⁵ Maranke Wieringa, *What to Account for When Accounting for Algorithms*, Proceedings of the 2020 ACM FAT Conference (Jan. 2020), <https://doi.org/10.1145/3351095.3372833>.

⁷⁶ Raji, Closing the AI Accountability Gap.

⁷⁷ Brundage, Toward Trustworthy AI Development.

Technical Glossary to the Key Considerations Appendix

This glossary provides a working set of definitions specific to the NSCAI Key Considerations. The Commission acknowledges that the definitions of the terms below may diverge from other scholarly or government definitions and were developed to be accessible to a broad audience.

AI Component: A software object that uses AI, meant to interact with other components, encapsulating certain functionality or a set of functionalities. An AI component has a clearly defined interface and conforms to a prescribed behavior common to all components within an architecture.¹

AI Lifecycle: The steps for managing the lifespan of an AI system: 1) Specify the system's objective. 2) Build model. 3) Test the AI system. 4) Deploy and maintain the AI system. 5) Engage in a feedback loop with continuous training and updates.²

AI System: A system designed or adapted to interact with an anticipated operational environment to achieve one or more intended purposes while complying with applicable constraints and that uses AI to provide a substantial part of its capabilities.³

Artificial Intelligence (AI): The ability of a computer system to solve problems and to perform tasks that have traditionally required human intelligence to solve.

Auditability: A characteristic of an AI system in which its software and documentation can be interrogated and yield information at each stage of the AI lifecycle to determine compliance with policy, standards, or regulations.

DevSecOps: Enhanced engineering practices that improve the lead time and frequency of delivery outcomes, promoting a more cohesive collaboration between development, security, and operations teams as they work toward continuous integration and delivery.

Differential Privacy: A criterion for a strong, mathematical definition of privacy in the context of statistical and ML analysis used to enable the collection, analysis, and sharing of a broad range of statistical estimates, such as averages, contingency tables, and synthetic data, based on personal data while protecting the privacy of the individuals in the data.⁴

False Negative: An example in which the predictive model mistakenly classifies an item as in the negative class. For example, a false negative describes the situation in which a junk-email model specifies that a particular email message is not spam (the negative class), when the email message actually is spam, leading to frustration of the junk message appearing in an end user's inbox.⁵ In a higher-stakes example, a false negative captures the case in which a medical diagnostic model misses identifying a disease that is present in a patient.

False Positive: An example in which the model mistakenly classifies an item as in the positive class. For example, the model inferred that a particular email message was spam (the positive class), but that email message was actually not spam, leading to delays in an end user reading a potentially important message.⁶ In a higher-stakes situation, a false positive describes the situation in which a disease is diagnosed as present when the disease is not present, potentially leading to unnecessary and costly treatments.

High-Fidelity Performance Traces: A commonly used technique useful in debugging and performance analysis. Concretely, trace recording implies detection and storage of relevant events during run-time, for later off-line analysis. High fidelity traces refers to the amount of fine-grained detail captured in the traces.⁷

Human Factors Engineering: The discipline that takes into account human strengths and limitations in the design of interactive systems that involve people, tools and technology, and work environments to ensure safety, effectiveness, and ease of use.⁸

Human in the Loop: The term describes a system architecture in which active human judgment and engagement are part of the operation of a system, and a human is an integral part of the system behavior. An example is the human operator of a remotely piloted vehicle or a decision support system that makes recommendations for a human to decide on.

Human on the Loop: This term describes a system architecture in which a human has a supervisory role in the operation of the system but is not an integral part of the system behavior. An example is an operator monitoring a fleet of warehouse robots—they operate autonomously but can be shut down if the operator determines something is wrong.

Machine Learning (ML): The study or the application of computer algorithms that improve automatically through experience.⁹ Machine learning algorithms build a model based on training data in order to perform a specific task, like aiding in prediction or decision-making processes, without necessarily being explicitly programmed to do so.

Model Testing: Testing assesses the performance of a trained model against new, previously unseen inputs, to demonstrate that the model generalizes to produce accurate results beyond just the training data.¹⁰

Model Training: Training a model simply means learning (determining) good values for all of the internal parameters that determine the model's performance. In supervised learning, for example, a machine learning model is trained by examining many labeled examples and attempting to find a model that minimizes the discrepancies between the real (labelled) values and the values produced by the model.¹¹

Technical Glossary to the Key Considerations Appendix

Multi-Party Federated Learning: A machine learning architecture in which many clients (e.g., mobile devices or whole organizations) collaboratively train a model under the orchestration of a central server (e.g., a service provider) while keeping the training data decentralized. It can mitigate many of the systemic privacy risks and costs resulting from traditional, centralized machine learning and data science approaches.¹² However, it does introduce new attack vectors that must be addressed.¹³

Precision: A metric for classification models. Precision identifies the frequency with which a model was correct when classifying the positive class. It answers the question “How many selected positive items are true positive?” For example, the percentage of messages flagged as spam that are spam.¹⁴

Privacy-Preserving AI: Techniques for protecting the privacy of people associated with the training data from adversarial attacks. These techniques include federated learning and differential privacy.¹⁵

Recall: A metric for classification models. Recall identifies the frequency with which a model correctly classifies the true positive items. It answers the question “How many true positive items were correctly classified?” For example, the percentage of spam messages that were flagged as spam.¹⁶

Reliable AI: An AI system that performs in its intended manner within the intended domain of use.

Robust AI: An AI system that is resilient in real-world settings, such as an object-recognition application that is robust to significant changes in lighting. The phrase also refers to resilience when it comes to adversarial attacks on AI components.

Run-Time Behavior: The behavior of a program while it is executing (i.e., running on one or more processors).

Trustworthy AI: Trustworthy AI has three components: (1) it should be lawful, ensuring compliance with all applicable laws and regulations; (2) it should be ethical, demonstrating respect for, and ensuring adherence to, ethical principles and values; and (3) it should be robust, both from a technical and social perspective, because, even with good intentions, AI systems can cause unintentional harm.¹⁷

Technical Glossary to the Key Considerations Appendix - Endnotes

¹ See NIST, *NISTIR 7298 Rev. 3, Glossary of Key Information Security Terms* (July 2019), <https://csrc.nist.gov/glossary/term/component>.

² Note that for data-driven AI systems step 2 is expanded and replaced with 2.a) Acquire data to meet the objective, and 2.b) Train the AI system on the data; and these two steps are usually repeated, with data acquisition and training continuing until desired performance objectives are attained. For further discussion on the ML lifecycle, see Saleema Amershi, et al., *Software Engineering for Machine Learning: A Case Study*, IEEE Computer Society (May 2019), <https://www.microsoft.com/en-us/research/publication/software-engineering-for-machine-learning-a-case-study/>.

³ See Hilary Sillitto, et al., *Systems Engineering and System Definitions*, International Council on Systems Engineering, (Jan. 8, 2019), https://www.incose.org/docs/default-source/default-document-library/final_se-definition.pdf.

⁴ Kobbi Nissim, et al., *Differential Privacy: A Primer for a Non-technical Audience*, Working Group of the Privacy Tools for Sharing Research Data Project, Harvard University, (Feb. 14, 2018), https://privacytools.seas.harvard.edu/files/privacytools/files/pedagogical-document-dp_new.pdf.

⁵ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

⁶ *Id.*

⁷ See Johan Kraft, et al., *Trace Recording for Embedded Systems: Lessons Learned from Five Industrial Projects*, Runtime Verification at 315-329, https://link.springer.com/chapter/10.1007%2F978-3-642-16612-9_24.

⁸ See *Human Factors Engineering*, U.S. Department of Health and Human Services: Agency for Healthcare Research and Quality (Sept. 2019), <https://psnet.ahrq.gov/primer/human-factors-engineering>.

⁹ Thomas M. Mitchell, *Machine Learning*, McGraw-Hill (1997).

¹⁰ See Rob Ashmore, et al., *Assuring the Machine Learning Lifecycle: Desiderata, Methods, and Challenges*, arXiv at 4 (May 2019), <https://arxiv.org/abs/1905.04223>.

¹¹ See *Descending into ML: Training and Loss*, Google (last accessed Feb. 15, 2021), <https://developers.google.com/machine-learning/crash-course/descending-into-ml/training-and-loss>.

¹² Peter Kairouz, et al., *Advances and Open Problems in Federated Learning*, arXiv (Dec. 10, 2019), <https://arxiv.org/pdf/1912.04977.pdf>.

¹³ See Vale Tolpegin, et al., *Data Poisoning Attacks Against Federated Learning Systems*, ArXiv (Aug. 11, 2020), <https://arxiv.org/abs/2007.08432>; Arjun Nitin Bhagoji, et al., *Analyzing Federated Learning Through an Adversarial Lens*, arXiv (Nov. 25, 2019), <https://arxiv.org/abs/1811.12470>.

¹⁴ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

¹⁵ For a discussion on how privacy-preserving machine learning works, see Roxanne Heston & Helon Toner, *Have Your Data and Use It Too: A Federal Initiative for Protecting Privacy While Advancing AI*, Day One Project (Jan. 23, 2020), <https://www.dayoneproject.org/post/have-your-data-and-use-it-too-a-federal-initiative-for-protecting-privacy-while-advancing-ai>; see also Georgios Kaissis, et al., *Secure, Privacy-Preserving and Federated Machine Learning in Medical Imaging*, Nature Machine Intelligence at 305-311 (June 8, 2020), <https://doi.org/10.1038/s42256-020-0186-1>.

¹⁶ See Frank Liang, *Evaluating the Performance of Machine Learning Models*, Towards Data Science (April 18, 2020), <https://towardsdatascience.com/classifying-model-outcomes-true-false-positives-negatives-177c1e702810>.

¹⁷ See *Ethics Guidelines for Trustworthy AI*, European Commission: High-Level Expert Group on Artificial Intelligence at 5 (April 8, 2019), <https://ec.europa.eu/futurium/en/ai-alliance-consultation>.

Appendix D:

Draft Legislative Language

The following legislative text represents the Commission staff's best efforts to capture the Commission's final recommendations in legislative form. The Commission defers to the House and Senate members, staff, and legislative counsels as to appropriate drafting.

CHAPTER 1: EMERGING THREATS IN THE AI ERA

Blueprint for Action

Combatting Malign Information Operations Enabled by AI

Recommendation: A National Strategy for the Global Information Domain.

Congress should direct the Executive Branch to transmit a National Strategy for the Global Information Domain that categorizes the global information domain as an arena of competition vital to the national security of the United States.

SEC. ____.—NATIONAL STRATEGY FOR THE GLOBAL INFORMATION DOMAIN.—

(a) IN GENERAL.—Not later than 270 days after the date of the enactment of this Act, the President shall transmit to Congress a National Strategy for the Global Information Domain that addresses the global information domain as an arena of competition vital to the national security of the United States.

(b) ISSUES ADDRESSED.—The National Strategy for the Global Information Domain required by subsection (a) shall, at a minimum:

(1) Prioritize the global information domain as an arena for international competition;

(2) Detail how adversarial state and non-state actors are attempting to define and control the global information domain in order to shape global opinion and achieve strategic advantage;

(3) Account for the critical role of artificial intelligence-enabled malign information in the efforts of adversarial state and non-state actors to achieve these goals;

(4) Identify and prioritize actions to defend, counter, and compete against malign information operations as a national security threat;

(5) As necessary, update critical infrastructure designations and require relevant departments and agencies to update sector-specific plans to reflect emerging technologies; and

(6) Establish organizational structures for U.S. national security agencies to counter and compete against the threat.

CHAPTER 2: FOUNDATIONS OF FUTURE DEFENSE

Blueprint for Action

Recommendation: Drive Change through Top-Down Leadership.

In the National Defense Authorization Act (NDAA) for Fiscal Year (FY) 2022, establish a Steering Committee on Emerging Technology and National Security Threats and designate that it be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.

SEC. ____.—ROLE OF INTELLIGENCE COMMUNITY IN STEERING COMMITTEE ON EMERGING TECHNOLOGY.—

Section 236 of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, is amended—

(1) in subsection (b), by—

(A) redesignating paragraph (8) as paragraph (9); and

(B) inserting the following new paragraph before redesignated paragraph (9):

“(8) One or more representatives of the Intelligence Community, to include the Principal Deputy Director of National Intelligence.”

(2) by redesignating paragraph (c) as paragraph (d); and inserting the following new paragraph before redesignated paragraph (d):

“(c) LEADERSHIP.—The Steering Committee shall be chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.”

The Steering Committee on Emerging Technology recommendation is also featured in Chapters 3 and 5.

Recommendation: Build the Technical Backbone.

Prioritize funding for the Department’s digital ecosystem and associated activities. The Armed Services Committees should use the FY 2022 NDAA to direct the Department of

Defense to develop a resourcing plan for the digital ecosystem that establishes, sustains, and incentivizes use of its various components as enterprise-wide, enduring resources. The Committees should also authorize the obligation of funds to begin work on the ecosystem.

SEC. ____.—RESOURCING PLAN FOR DIGITAL ECOSYSTEM.—

(a) IN GENERAL.—Within one year after the date of the enactment of this Act, the Secretary of Defense shall develop a plan for the development of a modern digital ecosystem that embraces state of the art tools and modern processes to enable development, testing, fielding, and continuous update of artificial intelligence-powered applications at speed and scale from headquarters to the tactical edge.

(b) CONTENTS OF PLAN.—At a minimum, the plan required by subsection (a) shall include—

(1) an open architecture and an evolving reference design and guidance for needed technical investments in the proposed ecosystem that address issues including common interfaces, authentication, applications, platforms, software, hardware, and data infrastructure; and

(2) a governance structure, together with associated policies and guidance, to drive the implementation of the reference throughout the Department on a federated basis.

Recommendation: Train and Educate Warfighters.

Component 1: Integrate Digital Skill Sets and Computational Thinking into Military Junior Leader Education.

Require the military services to integrate digital skills and computational thinking into pre-commissioning and entry-level training.

SEC. ____.—INTEGRATING DIGITAL SKILL SETS AND COMPUTATIONAL THINKING INTO MILITARY JUNIOR LEADER EDUCATION.—

Not later than 270 days after the date of the enactment of this Act, the Chief of Staff of the Army, the Chief of Naval Operations, the Chief of Staff of the Air Force, and the Commandant of the Marine Corps (collectively, the Service Chiefs) shall expand the curriculum for military junior leader education to incorporate appropriate training material related to problem definition and curation, a conceptual understanding of the artificial intelligence lifecycle, data collection and management, probabilistic reasoning and data visualization, and data-informed decision-making. Whenever possible, the new training and education should include the use of existing artificial intelligence-enabled systems and tools.

Component 2: Integrate Emerging and Disruptive Technologies into Service-level Professional Military Education.

Require the military services to integrate emerging and disruptive technologies into service-level Professional Military Education.

SEC. ____.—INTEGRATION OF MATERIAL ON EMERGING TECHNOLOGIES INTO PROFESSIONAL MILITARY EDUCATION.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense, in consultation with the Joint Chiefs of Staff, shall ensure that the curriculum for professional military education is revised in each of the military services to incorporate periodic courses on militarily significant emerging technologies that increasingly build the knowledge base, vocabulary, and skills necessary to intelligently analyze and utilize emerging technologies in the tactical, operational, and strategic levels of warfighting and warfighting support.

SEC. ____.—SHORT COURSE ON EMERGING TECHNOLOGIES FOR SENIOR CIVILIAN AND MILITARY LEADERS.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall establish a short course on emerging technologies for general and flag officers and senior executive-level civilian leaders. The short course shall be taught on an iterative, two-year cycle and shall address the most recent, most relevant technologies and how these technologies may be applied to military and business outcomes in the Department of Defense.

(b) THROUGHPUT OBJECTIVES.—In assessing participation in the short course authorized by subsection (a), the Secretary of Defense shall ensure that:

(1) In the first year that the course is offered, no fewer than twenty percent of general flag officers and senior executive-level civilian leaders are certified as having passed the short course required by subsection (a); and

(2) In each subsequent year, an additional ten percent of general flag officers and senior executive-level civilian leaders are certified as having passed such course, until such time as eighty percent of such officers and leaders are so certified.

Component 3: Create Emerging and Disruptive Technology Coded Billets in the Department of Defense.

Require the Department of Defense to create emerging and disruptive technology critical billets that must be filled by emerging technology certified leaders.

SEC. ____.—EMERGING TECHNOLOGY-CODED BILLETS WITHIN THE DEPARTMENT OF DEFENSE.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall ensure that the military services—

(1) code appropriate billets to be filled by emerging technology-qualified officers; and

(2) develop a process for officers to become emerging technology-qualified.

(b) APPROPRIATE POSITIONS.—Emerging technology-coded positions may include, as appropriate—

(1) positions responsible for assisting with acquisition of emerging technologies;

(2) positions responsible for helping integrate technology into field units;

(3) positions responsible for developing organizational and operational concepts;

(4) positions responsible for developing training and education plans; and

(5) leadership positions at the operational and tactical levels within the military services.

(c) QUALIFICATION PROCESS.—The process for qualifying officers for emerging technology-coded billets shall be modeled on a streamlined version of the joint qualification process and may include credit for serving in emerging technology focused fellowships, emerging technology focused talent exchanges, emerging technology focused positions within government, and educational courses focused on emerging technologies.

Recommendation: Accelerate Adoption of Existing Digital Technologies.

Component 3: Expand Use of Specialized Acquisition Pathways and Contracting Approaches.

Authorize the use of a rapid contracting mechanism for the software acquisition pathway.

SEC. ____.—RAPID CONTRACTING MECHANISM FOR SOFTWARE ACQUISITION.—

(a) IN GENERAL.—Not later than 270 days after the date of the enactment of this Act, the Secretary of Defense shall establish an agile contracting mechanism to support the software acquisition pathway developed pursuant to section 800 of the National Defense Authorization Act for Fiscal Year 2020 and embedded in Department of Defense Directives 5000.02 and 5000.87.

(b) CHARACTERISTICS.—The agile contracting mechanism established pursuant to subsection (a) shall authorize processes pursuant to which—

(1) a contract is awarded on the basis of statements of qualifications and past performance data submitted by contractors, supplemented by discussions with two or more contractors determined to be the most highly-qualified, without regard to price;

(2) the contract identifies the contractor team to be engaged for the work, and substitutions shall not be made during the base contract period without the advance written consent of the contracting officer;

(3) the contractor reviews existing software in consultation with the user community and utilizes user feedback to define and prioritize software requirements, and to design and implement new software and software upgrades, as appropriate;

(4) an independent, non-advocate cost estimate is developed in parallel with engineering of the software, leveraging agile cost estimation best practices rather than counting source lines of code; and

(5) value-based performance metrics are established and can be automatically generated by users to address issues such as deployment rate and speed of delivery, response rate such as the speed of recovery from outages and cybersecurity vulnerabilities, and assessment and estimation of the size and complexity of software development effort.

Component 4: Modernize the Budget and Oversight Processes for Digital Technologies.
Update title 10, Section 181 to designate USD(R&E) Co-Chair and Chief Science Advisor to the JROC.

SEC. ____.—ENHANCED ROLE OF UNDER SECRETARY OF DEFENSE FOR RESEARCH AND ENGINEERING ON THE JOINT REQUIREMENTS OVERSIGHT COUNCIL.—Section 181 of title 10, United States Code, is amended—

(1) in subsection (b), by—

(A) inserting “the Secretary of Defense and” before “the Chairman of the Joint Chiefs of Staff”;

(B) redesignating paragraphs (2) through (6) as paragraphs (3) through (7);

(C) inserting a new paragraph (2), as follows:

“(2) leveraging awareness of global technology trends, threats, and adversary capabilities to address gaps in joint military capabilities and validate technical feasibility of requirements developed by the military services;” and

(D) in redesignated paragraphs (4)(B) and (5) by inserting “the Secretary of Defense and” before “the Chairman of the Joint Chiefs of Staff”;

(2) in subsection (c), by—

(A) striking “Chairman of the Joint Chiefs of Staff for making recommendations about” in paragraph (1)(A) and inserting “Council for”;

(B) redesignating subparagraphs (B) through (E) of paragraph (1) as subparagraphs (C) through (F);

(C) adding a new paragraph (1)(B), as follows:

“(B) The Under Secretary of Defense for Research and Engineering, who is the co-Chair of the Council and is the Chief Science Advisor to the Council.”;

(D) by striking in paragraph (2) “(B), (C), (D), and (E)” and inserting “(C), (D), (E), and (F)”; and

(E) by amending paragraph (3) to read as follows:

“(3) In making any recommendation to the Secretary and the Chairman of the Joint Chiefs of Staff pursuant to this section, the Co-Chairs of the Council shall provide any dissenting view of members of the Council with respect to such recommendation.”; and

(3) in subsection (d), by—

(A) striking subparagraph (1)(D); and

(B) redesignating subparagraphs (E) through (H) of paragraph (1) as paragraphs (D) through (G).

Direct the Secretary of Defense to establish the dedicated AI fund.

SEC. ____.—ARTIFICIAL INTELLIGENCE DEVELOPMENT AND PROTOTYPING FUND.—

(a) IN GENERAL.—The Secretary of Defense shall establish a fund to be known as the “Artificial Intelligence Development and Prototyping Fund” to support operational prototyping and speed the transition of artificial intelligence-enabled applications into both service-specific and joint mission capabilities with priority on joint mission capabilities for Combatant Commanders. The Fund shall be managed by the Under Secretary of Defense for Research and Engineering, in consultation with the Joint Artificial Intelligence Center, the Joint Staff, and the military services.

(b) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to a military department for the purpose of carrying out a development or prototyping program selected by the Under Secretary of Defense for Research and Engineering for the purposes described in paragraph (1). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Department of Defense.

(c) CONGRESSIONAL NOTICE.—The Under Secretary of Defense for Research and Development shall notify the congressional defense committees of all transfers under paragraph (2). Each notification shall specify the amount transferred, the purpose of the transfer, and the total projected cost and estimated cost to complete the acquisition program to which the funds were transferred.

CHAPTER 3: AI AND WARFARE

Blueprint for Action

Recommendation: Establish AI-readiness performance goals.

Require the Secretary of Defense to establish performance objectives and accompanying metrics for AI and digital readiness and provide an update to Congress no later than 120 days after approving these goals.

SEC. ____.—ARTIFICIAL INTELLIGENCE READINESS GOALS.—

(a) IN GENERAL.—Not later than one year after the date of the enactment of this Act, the Secretary of Defense shall review the potential applications of artificial intelligence and digital technology to Department of Defense platforms, processes and operations, and establish performance objectives and accompanying metrics for the incorporation of artificial intelligence and digital readiness into such platforms, processes and operations.

(b) SKILLS GAPS.—As a part of the review required by subsection (a), the Secretary shall direct the military departments and defense components to—

(1) conduct a comprehensive review of skill gaps in the fields of software development, software engineering, knowledge management, data science, and artificial intelligence;

(2) assess the number and qualifications of civilian personnel needed for both management and specialist tracks in such fields;

(3) assess the number of military personnel (officer and enlisted) needed for both management and specialist tracks in such fields; and

(4) establish recruiting, training, and talent management goals to achieve and maintain staffing levels needed to fill identified gaps and meet the Department's needs for skilled personnel.

(c) REPORT TO CONGRESS.—Not later than 120 days after the completion of the review required by subsection (a), the Secretary shall report to Congress on the findings of the review and any action taken or proposed to be taken by the Secretary to address such findings.

Recommendation: Promote AI interoperability and the adoption of critical emerging technologies among allies and partners.

Component 6: Modify authorities and processes in order to improve DoD's ability to conduct international capability development.

SEC. ____.—ENHANCED AUTHORITY TO ENTER INTO COOPERATIVE RESEARCH AND DEVELOPMENT AGREEMENTS WITH INTERNATIONAL PARTNERS.—

(a) AUTHORITY OF SECRETARY OF DEFENSE.—Section 2350a of title 10, United States Code, is amended—

(1) In subsection (a), by—

(A) Adding a new subparagraph (F) at the end of paragraph (2), as follows:

“(F) Any business, academic or research institution, or other non-governmental entity organized pursuant to the laws of a country referred to in subparagraphs (C), (D) and (E), subject to the consent of the country involved.”;

(B) Amending paragraph (3) by striking “a country referred to in subparagraph (E) of paragraph (2),” and inserting “a country referred to in subparagraph (E) of paragraph (2) or a non-governmental entity referred to in subparagraph (F) of such paragraph,”; and

(C) Adding a new paragraph (4), as follows:

“(4) The Secretary may delegate the authority to enter memoranda of understanding pursuant to this section to the secretary of a military department, the Director of the Joint Artificial Intelligence Center, and the Director of the Defense Advanced Research Projects Agency, subject to such terms and conditions as may be necessary to ensure that any agreements entered are consistent with the foreign policy and defense policy of the United States.”; and

(2) In paragraph (1) of subsection (b), by striking “will improve, through the application of emerging technology,” and inserting “is likely to improve, through the application or enhancement of emerging technology,”;

(3) In subsection (c), by adding at the end the following new sentence:
 “If a foreign partner is expected to contribute significantly to the development of a new or novel capability, full consideration shall be given to non-monetary contributions, including the value of research and development capabilities and the strategic partnerships.”

(b) AUTHORITY OF THE PRESIDENT.—Section 2767 of title 22, United States Code, is amended—

(1) in subsection (c), by adding at the end the following new sentence:
 “If a foreign partner is expected to contribute significantly to the development of a new or novel capability, full consideration shall be given to non-monetary contributions, including the value of research and development capabilities and the strategic partnerships.”

(2) in subsection (f), by inserting before the semicolon in subparagraph (4) the following: “(and a description of any non-monetary contributions made by such participants)”; and

(3) in subsection (j), by—

(A) amending the title to read as follows: “Cooperative project agreements with friendly foreign countries not members of NATO and with non-governmental organizations in NATO and friendly non-NATO countries”; and

(B) amending paragraph (2) to read as follows:

“(2) The President may enter into a cooperative project agreement with any business, academic or research institution, or other non-governmental entity organized pursuant to the laws of NATO member or a friendly foreign country that is not a member of NATO, subject to the consent of the country involved.”

CHAPTER 5: AI AND THE FUTURE OF NATIONAL INTELLIGENCE

Blueprint for Action

Recommendation: Empower the IC’s science and technology leadership.

Designate the Director of S&T within ODNI as the IC CTO and grant that position additional authorities for establishing policies on, and supervising, IC research and engineering, technology development, technology transition, appropriate prototyping activities, experimentation, and developmental testing activities.

Grant the Director of National Intelligence sufficient budgetary authorities to enforce technical standards across the IC, including the ability to fence or otherwise withhold funding for programs that are not compliant with established common standards and policies.

SEC. ____.—CHIEF TECHNOLOGY OFFICER FOR THE INTELLIGENCE COMMUNITY.—

Section 3030 of title 50, United States Code, is amended—

(1) in subsection (a), by striking “who shall be appointed by the Director of National Intelligence” and inserting “who shall be appointed by the Director of National Intelligence and shall serve as the Chief Technology Officer for the Intelligence Community.”; and

(2) in subsection (c), by—

(A) redesignating paragraphs (2) through (5) as paragraphs (4) through (7); and

(B) inserting new paragraphs (2) and (3), as follows:

“(2) establish policies for the intelligence community on research and engineering, technology development, technology transition, prototyping activities, experimentation, and developmental testing, and oversee the implementation of such policies;

“(3) establish common technical standards and policies necessary to rapidly scale artificial intelligence-enabled applications across the intelligence community;”.

Suggested Report Language: The Chief Technology Officer for the Intelligence Community shall collect information on each Intelligence Community element’s compliance with applicable standards and policies for artificial intelligence research and development, and shall provide such information to the Director of National Intelligence. The Intelligence Committees encourage the Director of National Intelligence to closely review the compliance information and place a temporary hold on an Intelligence Community element that fails to execute artificial intelligence research and development funds in accordance with the applicable standards and policies.

Establish a fund that would allow the DNI to identify and invest in AI applications with outsized potential that may not have an identified source of agency or program funding as they near the end of their S&T life cycle.

SEC. ____.—ARTIFICIAL INTELLIGENCE CRITICAL APPLICATIONS FUND FOR THE INTELLIGENCE COMMUNITY.—

(a) IN GENERAL.—The Director of National Intelligence shall establish a fund

to be known as the “Artificial Intelligence Critical Applications Fund” to support agile development and fielding of artificial intelligence-enabled applications with exceptional potential for the intelligence community. The Fund shall be managed by the Director of Science and Technology, in consultation with the National Intelligence Science and Technology Committee established pursuant to section 3030 of title 50, United States Code.

(b) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to any element of the intelligence community for the purpose of carrying out a development or fielding program selected by the Director of Science and Technology for the purposes described in subsection (a). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Director of National Intelligence and the intelligence community.

(c) CONGRESSIONAL NOTICE.—The Director of National Intelligence shall notify the congressional intelligence committees and the congressional appropriations committees of all transfers under paragraph (2). Each notification shall specify the amount transferred, the purpose of the transfer, and the total projected cost and estimated cost to complete the acquisition program to which the funds were transferred.

Establish a 10-year, \$1,000,000,000 Program of Record to provide long-term, predictable funding for technologies identified in the technology annex to the National Intelligence Strategy.

SEC. ____.—ARTIFICIAL INTELLIGENCE TECHNOLOGY ROADMAP AND FUNDING PLAN FOR THE INTELLIGENCE COMMUNITY.—

(a) IN GENERAL.—Not later than 180 days after the date of the enactment of this Act, the Director of National Intelligence, in consultation with the Secretary of Defense, shall develop a technology annex to the National Intelligence Strategy and a ten-year plan to provide long-term, predictable funding of up to one billion dollars to implement the steps identified in such annex.

(b) CONTENTS OF TECHNOLOGY ANNEX.—The technology annex required by subsection (a) shall provide a technology roadmap for the adoption of artificial intelligence-enabled applications to solve operational intelligence requirements, including:

(1) A description of challenges faced in the intelligence community’s efforts to analyze the global environment and monitor technological advancements, adversarial capability development, and emerging threats;

(2) Identification of technical capabilities, including artificial intelligence capabilities, needed to enable steps to address each challenge;

(3) A prioritized, time-phased plan for developing or acquiring such technical capabilities, that takes into account research and development timelines, a strategy for public private partnerships, and a strategy for connecting researchers to end users for early prototyping, experimentation, and iteration;

(4) Any additional or revised acquisition policies and workforce training requirements that may be needed to enable intelligence community personnel to identify, procure, integrate, and operate the technologies identified in the annex;

(5) Identification of infrastructure requirements for developing and deploying technical capabilities, including:

(A) data, compute, storage, and network needs;

(B) a resourced and prioritized plan for establishing such infrastructure; and

(C) an analysis of the testing, evaluation, verification, and validation requirements to support prototyping and experimentation and a resourced plan to implement them, including standards, testbeds, and red-teams for testing artificial intelligence systems against digital “denial & deception” attacks.

(6) Consideration of human factor elements associated with priority technical capabilities, including innovative human-centric approaches to user interface, human-machine teaming, and workflow integration;

(7) Consideration of interoperability with allies and partners, including areas for sharing of data, tools, and intelligence products; and

(8) Flexibility to adapt and iterate annex implementation at the speed of technological advancement.

Recommendation: Improve coordination between the IC and DoD.

Revise the National Defense Authorization Act for Fiscal Year 2021 (FY 2021 NDAA) provision authorizing a Steering Committee on Emerging Technology by designating it to be tri-chaired by the Deputy Secretary of Defense, the Vice Chairman of the Joint Chiefs of Staff, and the Principal Deputy Director of National Intelligence.

See Chapter 2 recommendation “Drive Change through Top-Down Leadership” for proposed legislative text.

Recommendation: Aggressively pursue security clearance reform for clearances at the Top Secret level and above, and enforce security clearance reciprocity among members of the IC.

Congress should require the DNI to develop an implementation plan for security clearance reform for clearances at the Top Secret and above level including detailed timelines and metrics.

Congress should require the DNI and the directors of the major intelligence services to regularly report on progress to the oversight committees.

SEC. ____.—IMPLEMENTATION PLAN FOR SECURITY CLEARANCE REFORM.—

(a) PLAN REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Director of National Intelligence shall develop an implementation plan for security clearance reform for clearances at the Top Secret level and above. The implementation plan shall include, at a minimum:

(1) detailed implementation metrics and timelines;

(2) steps to be taken to collaborate with the private sector and academia to develop data-informed behavioral approaches to understanding risk factors and security clearance adjudication; and

(3) steps to be taken to reform identity management and ensure seamless security clearance reciprocity across the intelligence community (including any enforcement mechanisms that may be needed to ensure such reciprocity).

(b) REPORTS REQUIRED.—Not later than 270 days after the date of the enactment of this Act and annually for five years thereafter, the Director of National Intelligence shall report to the congressional intelligence committees on the implementation of the plan required by subsection (a) and the progress that has been made toward security clearance reform.

CHAPTER 6: TECHNICAL TALENT IN GOVERNMENT

Blueprint for Action

Recommendation: Create a National Reserve Digital Corps.

NATIONAL RESERVE DIGITAL CORPS ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “National Reserve Digital Corps Act of 2021”.

SEC. 2.—ESTABLISHMENT OF NATIONAL RESERVE DIGITAL CORPS.—

(a) IN GENERAL.—Subpart I of part III of title 5, United States Code, is amended by inserting after chapter 102 the following new chapter:

CHAPTER 103—NATIONAL RESERVE DIGITAL CORPS

SEC. 10301. Establishment.

SEC. 10302. Definitions.

SEC. 10303. Organization.

SEC. 10304. Work on Behalf of Federal Agencies.

SEC. 10305. Digital Corps Scholarship Program.

SEC. 10306. Duration of Pilot Program.

SEC. 10307. Authorization of Appropriation.

SEC. 10301. ESTABLISHMENT.—For the purposes of attracting, recruiting, and training a corps of world-class digital talent to serve the national interest and enable the Federal Government to become a digitally proficient enterprise, there is established within the Office of Management and Budget a pilot program for a civilian National Reserve Digital Corps, whose members shall serve as special government employees, working not fewer than 30 days per year as short-term advisors, instructors, or developers in the Federal Government.

SEC. 10302. DEFINITIONS.—

(a) DIRECTOR.—The term “Director” means the Director of the Office of Management and Budget.

(b) NODE.—The term “node” means a group of persons or team organized under the direction of a node leader to provide digital service to one or more Federal agencies pursuant to an agreement between the Office of Management Budget and each other Federal agency.

(c) NODE LEADER.—The term “node leader” means a full time government employee, as defined by section 2105 of title 5, United States Code, selected under this Act to lead one or more nodes, who reports to the Director or the Director’s designee.

(d) NODE MEMBER.—The term “node member” means a special government employee, as defined by section 202 of title 18, United States Code, selected under this Act to work at least 38 days per fiscal year and report to a node leader in furtherance of the mission of a specified node.

SEC. 10303. ORGANIZATION.—

(a) NODES AND NODE LEADERS.—The National Reserve Digital Corps shall be organized into nodes, each of which shall be under the supervision of a node leader .

(b) ADMINISTRATIVE SUPPORT.—The National Reserve Digital Corps shall receive funding and administrative support from the Office of Management and Budget,

which shall be responsible for selecting node leaders, establishing standards, ensuring that nodes meet government client requirements, maintaining security clearances, establishing access to an agile development environment and tools, and facilitating appropriate technical exchange meetings.

(c) **HIRING AUTHORITY.—**

(1) **Direct Hiring Authority of Node Members.—**The Director of the Office of Management and Budget, on the recommendation of a node leader, may appoint, without regard to the provisions of subchapter I of chapter 33 (other than sections 3303 and 3328 of such chapter), a qualified candidate to a position in the competitive service in the Office of Management and Budget to serve as a node member. This provision shall not preclude the Director from hiring additional employees, including full time government employees, as defined by section 2105 of title 5, United States Code.

(2) **Term and Temporary Appointments of Node Members.—**The Director of the Office of Management and Budget, on the recommendation of a node leader, may make a noncompetitive temporary appointment or term appointment for a period of not more than 18 months, of a qualified candidate to serve as a node member in a position in the competitive service for which a critical hiring need exists, as determined under section 3304 of title 5, United States Code, without regard to sections 3327 and 3330 of such title.

SEC. 10304. WORK ON BEHALF OF FEDERAL AGENCIES.—

(a) **PURPOSE.—**Each node shall undertake projects to assist Federal agencies by providing digital education and training, performing data triage and providing acquisition assistance, helping guide digital projects and frame technical solutions, helping build bridges between public needs and private sector capabilities, and related tasks.

(b) **AUTHORITIES.—**Projects may be undertaken—

(1) on behalf of a Federal agency—

(A) by direct agreement between the Office of Management and Budget and the Federal agency; or

(B) at the direction of the Office of Management and Budget at the request of the Federal agency; or

(2) to address a digital service need encompassing more than one Federal agency—

(A) at the direction of the Office of Management and Budget; or

(B) on the initiative of a node leader.

SEC. 10305. DIGITAL CORPS SCHOLARSHIP PROGRAM.—

(a) IN GENERAL.—The Director shall establish a National Reserve Digital Corps scholarship program to provide full scholarships to competitively selected students who commit to study specific disciplines related to national security digital technology .

(b) SERVICE OBLIGATION.—Each student, prior to commencing the Digital Corps Scholarship Program, shall sign an agreement with respect to the student's commitment to the United States. The agreement shall provide that the student agree to the following:

(1) a commitment to serve as an intern in a Federal agency for at least six weeks during each of the summers before their junior and senior years; and

(2) a commitment to serve in the National Reserve Digital Corps for six years after graduation.

(c) PROGRAM ELEMENTS.—In establishing the program, the Director shall determine the following—

(1) Eligibility standards for program participation;

(2) Criteria for establishing the dollar amount of a scholarship, including tuition, room and board;

(3) Repayment requirements for students who fail to complete their service obligation;

(4) An approach to ensuring that qualified graduates of the program are promptly hired and assigned to node leaders; and

(5) Resources required for the implementation of the program.

(d) CONTINUING EDUCATION.—The Director shall establish a training and continuing education program to fund educational opportunities for members of the National Digital Reserve Corps, including conferences, seminars, degree and certificate granting programs, and other training opportunities that are expected to increase the digital competencies of the participants.

(e) IMPLEMENTATION.—

(1) Not later than six months after the date of the enactment of this Act, the Director shall establish the administrative support function and issue guidance for the National Reserve Digital Corps, which shall include the identification of points of contact for node leaders at Federal agencies.

(2) Not later than one year after the date of the enactment of this Act, the Director shall appoint not fewer than five node leaders under the National Reserve Digital Corps program and authorize the node leaders to begin recruiting reservists and undertaking projects for Federal agencies.

(3) Beginning two years after the date of the enactment of this Act, the Director shall report annually to Congress on the progress of the National Reserve Digital Corps. The Director's report shall address, at a minimum, the following measures of success:

(A) The number of technologists who participate in the National Reserve Digital Corps annually;

(B) Identification of the Federal agencies that submitted work requests, the nature of the work requests, which work requests were assigned a node, and which work requests were completed or remain in progress;

(C) Evaluations of results of National Reserve Digital Corps projects by Federal agencies; and

(D) Evaluations of results of National Reserve Digital Corps projects by reservists.

SEC. 10306. DURATION OF PILOT PROGRAM.—The pilot program under this Act shall terminate no earlier than six years after its commencement.

SEC. 10307. AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated \$16,000,000 to remain available until fiscal year 2023 the initial administrative cost, including for the salaries and expenses scholarship and education benefits, for the National Digital Reserve Corps.

Recommendation: Create Digital Talent Recruiting Offices Aligned with Digital Corps.

SEC. ____.—DIGITAL TALENT RECRUITING OFFICES.—

(a) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF DEFENSE.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Defense shall designate a chief digital recruiting officer within the office of the Under Secretary of Defense for Personnel and Readiness to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Defense needs for specific types of digital talent;

(B) recruiting technologists, in partnership with the military services and defense components, including by attending conferences and career fairs, and actively recruiting on university campuses and from the private sector;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the military services and defense components to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Defense shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(b) DIGITAL TALENT RECRUITING FOR THE INTELLIGENCE COMMUNITY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Director of National Intelligence shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying intelligence community needs for specific types of digital talent;

(B) recruiting technologists, in partnership with components of the intelligence community, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into intelligence community recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the components of the intelligence community to use direct-hire authorities to accelerate hiring.

(3) The Director of National Intelligence shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(c) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF HOMELAND SECURITY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Homeland Security shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Homeland Security needs for specific types of digital talent;

(B) recruiting technologists, in partnership with components of the Department of Homeland Security, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in the components of the Department of Homeland Security to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Homeland Security shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

(d) DIGITAL TALENT RECRUITING FOR THE DEPARTMENT OF ENERGY.—

(1) Not later than 270 days after the date of the enactment of this Act, the Secretary of Energy shall designate a chief digital recruiting officer to oversee a digital recruiting office to carry out the responsibilities set forth in paragraph (2).

(2) The chief digital recruiting officer shall be responsible for—

(A) identifying Department of Energy needs for specific types of digital talent;

(B) recruiting technologists, in partnership with Department of Energy programs, by attending conferences and career fairs, and actively recruiting on college campuses;

(C) integrating Federal scholarship for service programs into civilian recruiting;

(D) offering recruitment and referral bonuses; and

(E) partnering with human resource teams in Department of Energy programs to use direct-hire authorities to accelerate hiring.

(3) The Secretary of Energy shall ensure that the chief digital recruiting officer is provided with personnel and resources sufficient to maintain an office and to carry out the duties set forth in paragraph (2).

Recommendation: Grant exemption from OPM General Schedule Qualification Policies for Specific Billets and Position Descriptions.

SEC. ____.—WAIVER OF QUALIFICATION STANDARDS FOR GENERAL SCHEDULE POSITIONS IN ARTIFICIAL INTELLIGENCE.—

(a) DEPARTMENT OF DEFENSE.—Two-star and above commands and their civilian equivalents are authorized to waive any General Schedule qualification standard established by the Office of Personnel Management in the case of any applicant for a position in artificial intelligence who is determined by a hiring manager, in consultation with subject matter experts, to be the best qualified candidate for the position.

(b) OTHER NATIONAL SECURITY AGENCIES.—The Director of the Office of Personnel Management shall establish a process by which the the Attorney General, the Secretary of Homeland Security, the Secretary of State, the Secretary of Commerce, the Director of National Intelligence, and the head of any element of the Intelligence Community may request an exception to any General Schedule qualification standard in any case in

which the agency head determines that national security needs would best be met by hiring managers making an independent judgment about qualifications and pay grades for a position in artificial intelligence with the advice of subject matter experts. The process shall provide for requests to be made for individual billets, for position descriptions, or for categories of individual billets or position descriptions at the discretion of the agency head.

Recommendation: Expand the CyberCorps: Scholarship for Service.

SEC. ____.—AMENDMENT TO THE FEDERAL CYBER SCHOLARSHIP-FOR- SERVICE PROGRAM.—

(a) AMENDMENTS TO TITLE 15, UNITED STATES CODE.—Section 7442 of title 15, United States Code, is amended—

(1) By amending the title to read: “Federal Cyber and Artificial Intelligence Scholarship-for-Service Program”;

(2) in subsection (a), by striking “industrial control system” and all that follows and inserting in lieu thereof “digital engineers, artificial intelligence practitioners, data engineers, data analysts, data scientists, industrial control system security professionals, security managers, and cybersecurity course instructors to meet the needs of the cybersecurity and artificial intelligence missions for Federal, State, local, tribal, and territorial governments.”;

(3) in subsection (b), by—

(A) striking “and” at the end of paragraph (3);

(B) striking the period at the end of paragraph (4) and inserting in lieu thereof “; and”; and

(C) adding a new paragraph (5), as follows:

“(5) provide an opportunity for scholarship recipients to initiate the security clearance process at least one year before their planned graduation date.”; and

(4) in subsection (c), by striking “3 years” and inserting “4 years”.

(b) SAVINGS PROVISION.—Nothing in this section, or an amendment made by this section, shall affect any agreement, scholarship, loan, or repayment under section 302 of the Cybersecurity Enhancement Act of 2014 (15 U.S.C. 7442), in effect on the day before the date of the enactment of this section.

Recommendation: Create a United States Digital Service Academy.

UNITED STATES DIGITAL SERVICE ACADEMY ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “United States Digital Service Academy Act of 2021”.

SEC. 2.—ESTABLISHMENT OF ACADEMY.—

(a) ESTABLISHMENT.—There is established as an independent entity within the Federal Government a United States Digital Service Academy (hereafter referred to as the “ACADEMY”), at a location to be determined, to serve as a federally-funded, accredited, degree-granting university for the instruction of selected individuals in digital technical fields and the preparation of selected individuals for civil service with the Federal Government.

(b) DIGITAL TECHNICAL FIELDS DEFINED.—The term “digital technical fields” includes artificial intelligence, software engineering, electrical science and engineering, computer science, molecular biology, computational biology, biological engineering, cybersecurity, data science, mathematics, physics, human-computer interaction, robotics, and design and any additional fields specified in regulations by the Board.

SEC. 3.—ORGANIZATION.—

(a) BOARD OF REGENTS.—The business of the Academy shall be conducted by a Board of Regents (hereafter referred to as the “Board”).

(1) COMPOSITION.—The Board shall consist of nine voting members and ex officio members, as set forth in this subsection.

(2) VOTING MEMBERS.—The President shall appoint, by and with the consent of the Senate, nine persons from civilian life who have demonstrated achievement in one or more digital technical fields, higher education administration, or Federal civilian service, to serve as voting members on the Board. Appointment of the first voting members shall be made not later than 180 days after enactment of this Act.

(3) EX OFFICIO MEMBERS.—Ex officio members shall include—

(A) The Secretary of State;

(B) The Secretary of Defense;

(C) The Attorney General;

(D) The Secretary of Commerce;

(E) The Secretary of Energy;

(F) The Secretary of Homeland Security;

(G) The Director of National Intelligence;

(H) The Director of the Office of Personnel Management; and

(I) such other Federal Government officials as determined by the President.

(2) TERM OF VOTING MEMBERS.—The term of office of each voting member of the Board shall be six years, except that initial terms shall be staggered at two year intervals and any member appointed to fill a vacancy occurring before the expiration of a term shall be appointed for the remainder of such term.

(3) PRESIDENT OF THE BOARD.—One of the members (other than an ex officio member) shall be designated by the President as Chairman and shall be the presiding officer of the Board.

(b) KEY POSITIONS.—There shall be at the Academy the following:

(1) A Superintendent;

(2) A Dean of the Academic Board, who is a permanent professor;

(3) A Director of Admissions; and

(4) A Director of Placement.

(c) SUPERINTENDENT.—The Board shall appoint a Superintendent of the Academy, who shall serve for a term of six years. The Superintendent, acting pursuant to the oversight and direction of the Board, shall be responsible for the day-to-day operations of the Academy and the welfare of the students and the staff of the Academy. The Board shall select the first Superintendent of the Academy no later than 60 days after the Board is established.

(d) ADVISORY BOARD.—The Board of Regents and the Superintendent shall be assisted by an Advisory Board, composed of commercial and academic leaders in digital technical fields and higher education. The Advisory Board shall adhere to the requirements of the Federal Advisory Committee Act, Pub.L. 92–463.

(e) INTERAGENCY WORKING GROUP.—

(1) ESTABLISHMENT.—The Office of Personnel Management shall establish and lead an interagency working group to annually assess and report to the Academy the need for civil servants at agencies in digital technical fields for the purposes of informing Academy student field of study and agency placement.

(2) RESPONSIBILITIES.—The interagency working group shall be responsible for—

(A) establishing a range of Academy graduates needed during the ensuing five-year period, by agency and digital technical field; and

(B) undertaking necessary steps to enable each agency identified to hire Academy graduates into full-time positions in the civil service.

(3) COMPOSITION.—The interagency working group shall consist of the following officials or their designees:

(A) The Secretary of State;

(B) The Secretary of Defense;

(C) The Attorney General;

(D) The Secretary of Commerce;

(E) The Secretary of Energy;

(F) The Secretary of Homeland Security;

(G) The Director of National Intelligence;

(H) The Director of the Office of Personnel Management; and

(I) such other Federal Government officials as determined by the Director of the Office of Personnel Management.

SEC. 4.—FACULTY.—

(a) NUMBER OF FACULTY.—The Superintendent of the Academy may employ as many professors, instructors, and lecturers at the Academy as the Superintendent considers necessary to achieve academic excellence.

(b) FACULTY COMPENSATION.—The Superintendent may prescribe the compensation of persons employed under this section. Compensation and benefits for faculty members of the Academy shall be sufficiently competitive to achieve academic excellence, as determined by the Superintendent.

(c) FACULTY EXPECTATIONS.—Faculty members shall—

(1) possess academic expertise and teaching prowess;

(2) exemplify high standards of conduct and performance;

(3) be expected to participate in the full spectrum of academy programs, including providing leadership for the curricular and extracurricular activities of students;

(4) comply with the standards of conduct and performance established by the Superintendent; and

(5) participate actively in the development of the students through the enforcement of standards of behavior and conduct, to be established in the Academy's rules and regulations.

(d) DEPARTMENT TITLES.—The Superintendent may prescribe the titles of each of the departments of instruction and the professors of the Academy.

SEC. 5.—STUDENT QUALIFICATIONS AND REQUIREMENTS FOR ADMISSION.—

(a) ADMISSIONS REQUIREMENTS.—A student wishing to be admitted to the Academy shall fulfill admission requirements to be determined by the Superintendent and approved by the Board of Regents.

(b) HONOR CODE.—A student wishing to be admitted to the Academy shall sign an Honor Code developed by the Superintendent of the Academy and approved by the Board of Regents. A violation of the honor code may constitute a basis for dismissal from the Academy.

SEC. 6.—APPOINTMENT OF STUDENTS.—

(a) NOMINATIONS PROCESS.—Prospective applicants to the Academy for seats described in paragraphs (1) and (2) of subsection (b) shall follow a nomination process established by the Director of Admissions of the Academy that is similar to the process used for admission to the military academies of the United States Armed Forces.

(b) APPOINTMENTS.—

(1) NOMINEES FOR CONGRESSIONAL SEATS.—Each member of the Senate or the House of Representatives may nominate candidates from the State that the member represents for each incoming first-year class of the Academy .

(2) EXECUTIVE BRANCH NOMINEES.—The President may nominate a maximum of 75 candidates to compete for the executive branch seats.

SEC. 7.—ACADEMIC FOCUS OF THE UNITED STATES DIGITAL SERVICE ACADEMY.—

(a) CURRICULUM.—Each Academy student shall follow a structured curriculum according to the program of study approved by the Board of Regents centered on digital technical fields and incorporating additional core curriculum coursework in history, government, English language arts including composition, and ethics.

(b) DEGREES CONFERRED UPON GRADUATION.—Under such conditions as the Board of Regents may prescribe, once the Academy is accredited, the Superintendent of the Academy may confer a baccalaureate of science or baccalaureate of arts degree upon a graduate of the Academy.

(c) MAJORS AND AREAS OF CONCENTRATION.—Under such conditions as the Board of Regents may prescribe, the Superintendent of the Academy may prescribe requirements for majors and concentrations and requirements for declaring a major or concentration during the course of study.

(d) ADDITIONAL DIGITAL SERVICE OF CIVIL SERVICE PROGRAMMING.—Under such conditions as the Board of Regents may prescribe, the Superintendent of the Academy may prescribe requirements for each Academy student to participate in non-curricular programming during Academy terms and during the summer, which may include internships, summer learning programs, and project-based learning activities.

SEC. 8.—CIVIL SERVICE REQUIREMENTS FOLLOWING GRADUATION.—

(a) CIVIL SERVICE AGREEMENT.—Each Academy student, prior to commencing the third year of coursework, shall sign an agreement with respect to the student's length of civil service to the United States. The agreement shall provide that the student agrees to the following:

(1) The student will complete the course of instruction at the Academy, culminating in graduation from the Academy.

(2) Unless the student pursues graduate education under subsection (f), upon graduation from the Academy, the student agrees to serve in the Federal civil service for not less than five years following graduation from the Academy .

(b) FAILURE TO GRADUATE.—

(1) IN GENERAL.—An Academy student who has completed a minimum of four semesters at the Academy but fails to fulfill the Academy's requirements for graduation shall be—

(A) dismissed from the Academy; and

(B) obligated to repay the Academy for the cost of the delinquent student's education in the amount described in paragraph (2).

(2) AMOUNT OF REPAYMENT.—A student who fails to graduate shall have financial responsibility for certain costs relating to each semester that the student was officially enrolled in the Academy as prescribed by the Superintendent.

(c) FAILURE TO ACCEPT OR COMPLETE ASSIGNED CIVIL SERVICE.—

(1) IN GENERAL.—A student who graduates from the Academy but fails to complete the full term of required civil service shall be obligated to repay the Academy for a portion of the cost of the graduate's education as determined by Academy as set forth in this subsection.

(2) AMOUNT OF REPAYMENT.—In the case of a delinquent graduate who fails to complete all years of public service required under subsection (a)(2) (including any additional years required for graduate education under subsection (f)), the delinquent graduate shall be financially responsible for the cost of the delinquent graduate's education (including the costs of any graduate education), except that the amount of financial responsibility under this paragraph shall be reduced by 20 percent for each year of civil service under subsection (a)(2) that the delinquent graduate did complete.

(d) EXCEPTIONS.—The Superintendent may provide for the partial or total waiver or suspension of any civil service or payment obligation by an individual under this section whenever compliance by the individual with the obligation is impossible or deemed to involve extreme hardship to the individual, or if enforcement of such obligation with respect to the individual would be unconscionable.

(e) STUDENT SALARIES AND BENEFITS.—The Academy shall not be responsible for the salaries and benefits of graduates of the Academy while the graduates are fulfilling the civilian service assignment under this section. All salaries and benefits shall be paid by the employer with whom the Academy graduate is placed.

(f) GRADUATE EDUCATIONS.—An Academy student and the Superintendent may modify the agreement under subsection (a) to provide that—

(1) the Academy shall—

(A) subsidize an Academy student's graduate education; and

(B) postpone the public service assignment required under subsection (a)(2).

(2) the student shall—

(A) accept a civil service assignment under subsection (g) upon the student's completion of the graduate program; and

(B) add two additional years to the student's civil service commitment required under the agreement described in subsection (a) for every year of subsidized graduate education.

SEC. 9.—IMPLEMENTATION PLAN.—

(a) Not later than 180 days after the enactment of this Act, the Superintendent, in consultation with the Advisory Board, shall develop a detailed plan to implement the Academy that complies with the requirements of this section. Upon approval by the Board of Regents, the Superintendent shall present the implementation plan to Congress.

(b) CONTENTS OF PLAN.—The implementation plan described in section (a) shall provide, a minimum, the following:

(1) Identification and securement of an appropriate site for initial Academy build-out with room for future expansion, to include a construction plan and temporary site plan, if necessary;

(2) Identification of gaps in the government's current and envisioned digital workforce by the interagency working group under the Office of Personnel Management as established by section (3)(e);

(3) Establishment of student qualifications and requirements for admission;

(4) Establishment of the student appointment and nomination process;

(5) Establishment of student honor and conduct code to include a plan for student noncompletion of requirements and obligations;

(6) Establishment of the student curriculum;

(7) Establishment of a mechanism for students to select fields of study and annually select agencies and career fields within the limits prescribed by

the interagency working group under the Office of Personnel Management as established by section (3)(e);

(8) Establishment of a mechanism for graduates to transition from the Academy to civil service employment by selected individual agencies;

(9) Determination of the initial Academy departments and faculty needs;

(10) Establishment of faculty and staff requirements and compensation;

(11) Determination of non-academic staff required;

(12) Recruitment and hiring of faculty, including tenure-track faculty, adjunct faculty, part-time faculty and visiting faculty, and other staff as needed;

(13) Identification of nonprofit and private sector partners;

(14) Procurement of outside funds and gifts from individuals and corporations for startup, administrative, maintenance, and infrastructure costs;

(15) Establishment of the process to meet statutory and regulatory requirements for establishing the Academy as an academic institution with degree-granting approval and for applying for degree program specific accreditation and ensuring that the Academy obtains, no later than two years after enactment of this Act, status as an accreditation candidate, as defined by a nationally recognized accrediting agency or association as determined by the Secretary of Education in accordance with section 1099b in title 10, United States Code, before commencing academic operations;

(16) A plan commencing the Academy with an initial class of 500 students three years after enactment of this Act;

(17) Procedures for incorporating accreditation assessments to facilitate ongoing improvements to the Academy; and,

(18) Procedures for assessing the size of the Academy and potential expansion of student enrollment.

SEC. 10.—ADMINISTRATIVE MATTERS.—

(a) FULLY-SUBSIDIZED EDUCATION.—Each Academy student's tuition and room and board shall be fully subsidized provided that the student completes the requirements of the Academy and fulfills the civil service commitment as determined by the implementation plan in section 9.

(b) GIFT AUTHORITY.—The Board of Regents may accept, hold, administer, and spend any gift, devise, or bequest of real property, personal property, or money made on the condition that the gift, devise, or bequest be used for the benefit, or in connection with, the establishment, operation, or maintenance, of the Academy. The Board of Regents may accept a gift of services, which includes activities that benefit the education, morale, welfare, or recreation of students, faculty or staff, for the Academy.

(1) LIMITATIONS AND PROHIBITIONS.—

(A) IN GENERAL.—The Board of Regents may not accept a gift under this subsection if the acceptance of the gift would reflect unfavorably on the ability of any agency of the Federal Government to carry out any responsibility or duty in a fair and objective manner, or would compromise the integrity or appearance of integrity of any program of the Federal Government or any officer or employee of the Federal Government who is involved in any such program.

(B) FOREIGN GIFTS.—The Board of Regents may not accept a gift of services from a foreign government or international organization under this subsection. A gift of real property, personal property, or money from a foreign government or international organization may be accepted under this subsection only if the gift is not designated for a specific individual.

(C) APPLICABLE LAW.—No gift under this section may be accepted with attached conditions inconsistent with applicable law or regulation.

(D) MISSION.—No gift under this section may be accepted with attached conditions inconsistent with the mission of the Academy .

(E) NAMING RIGHTS.—The Board of Regents may issue regulations governing the circumstances under which gifts conditioned on naming rights may be accepted, appropriate naming conventions, and suitable display standards.

(2) TREATMENT OF GIFTS.—

(A) Gifts and bequests of money, and the proceeds of the sale of property, received under subsection shall be deposited in the Treasury in the account of the Academy as no year money and may be expended in connection with the activities of the Academy as determined by the Board of Regents.

(B) The Board of Regents may pay all necessary expenses in connection with the conveyance or transfer of a gift, devise, or bequest accepted under this section.

(C) For the purposes of Federal income, estate, and gift taxes, any property, money, or services accepted under this subsection shall be considered as a gift, devise, or bequest to or for the use of the United States.

(D) The Comptroller General shall make periodic audits of gifts, devises, and bequests accepted under this section at such intervals as the Comptroller General determines to be warranted. The Comptroller General shall submit to Congress a report on the results of each such audit.

SEC. 11.—INITIAL APPROPRIATION.—There are authorized to be appropriated \$40,000,000 to remain available until expended for the Academy's initial administrative cost and salaries and expenses.

Recommendation: Establish Career Fields for Government Civilians in Software Development, Software Engineering, Data Science, Knowledge Management, and Artificial Intelligence.

SEC. ____.—NEW OCCUPATIONAL SERIES FOR DIGITAL CAREER FIELDS.—Not later than 270 days after the date of the enactment of this Act, the Director of the Office of Personnel Management shall exercise its authority under section 5105 of title 5, United States Code, to establish one or more new occupational series and associated policies covering Federal Government positions in the fields of software development, software engineering, data science, and knowledge management.

SEC. ____.—NEW OCCUPATIONAL SERIES FOR ARTIFICIAL INTELLIGENCE.—Not later than 270 days after the date of the enactment of this Act, the Director of the Office of Personnel Management shall exercise its authority under section 5105 of title 5, United States Code, to establish a new occupational series and associated policies covering Federal Government positions in the field of artificial intelligence.

Recommendation: Establish Digital Career Fields for Military Personnel.

SEC. ____.—MILITARY CAREER FIELDS FOR SOFTWARE DEVELOPMENT, DATA SCIENCE, AND ARTIFICIAL INTELLIGENCE.—Section 230 of the National Defense Authorization Act for Fiscal Year 2020 is amended by adding the following new subsection: “(d) Not later than 270 days after the date of the enactment of this subsection, the Chief of Staff of the Army, the Chief of Naval Operations, the Chief of Staff of the Air Force, and

the Commandant of the Marine Corps (collectively, the Service Chiefs) shall each establish new military career fields for software development, data science, and artificial intelligence that are open to commissioned officers, enlisted personnel and, as appropriate, warrant officers. The Service Chiefs shall utilize the authority provided in sections 605 and 649a to 649k of title 10, United States Code, to ensure that military personnel in these career fields who choose to specialize and focus on technical skill sets rather than pursue leadership positions are not required to move outside their specialties or into management positions to continue to promote.

CHAPTER 8: UPHOLDING DEMOCRATIC VALUES: PRIVACY, CIVIL LIBERTIES, AND CIVIL RIGHTS IN USES OF AI FOR NATIONAL SECURITY

Blueprint for Action

Recommendation Set 1: Increase Public Transparency about AI Use through Improved Reporting.

For AI systems that involve U.S. persons, require AI Risk Assessment Reports and AI Impact Assessments to assess the privacy, civil liberties and civil rights implications for each new qualifying AI system or significant system refresh.

SEC. ___—PRIVACY, CIVIL RIGHTS AND CIVIL LIBERTIES RISK AND IMPACT ASSESSMENTS FOR ARTIFICIAL INTELLIGENCE SYSTEMS.—

(a) IN GENERAL.—The head of a covered agency shall conduct risk and impact assessments of the privacy, civil rights, and civil liberties risks and potential implications of any covered artificial intelligence system utilized by the covered agency and take appropriate steps to mitigate risks and adverse impact of any such system on the privacy, civil rights, and civil liberties of U.S. persons.

(b) DEFINITIONS.—For purposes of this section—

(1) COVERED ARTIFICIAL INTELLIGENCE SYSTEM.—A “covered artificial intelligence system” means a qualified artificial intelligence system or a significant artificial intelligence system refresh as determined by the task force established in section [XX] of this Act that is—

(A) designed to collect, process, maintain, or use information on U.S. persons;

(B) may inadvertently process, maintain, or use information on U.S. persons; or

(C) has a direct impact on U.S. persons.

(2) COVERED AGENCY.—A “covered agency” includes—

(A) the Department of Homeland Security;

(B) the Federal Bureau of Investigation; and

(C) each element of the Intelligence Community, as defined in section 3003(4) of title 50, United States Code.

(3) HEAD OF A COVERED AGENCY.—The “head of a covered agency” shall mean the Secretary of Homeland Security, the Director of the Federal Bureau of Investigation and, for the Intelligence Community, the Director of National Intelligence.

(c) REPORTS REQUIRED.—

(1) ARTIFICIAL INTELLIGENCE SYSTEM RISK ASSESSMENT.—Before acquiring or fielding a covered artificial intelligence system, each covered agency shall conduct an Artificial Intelligence System Risk Assessment (“Risk Assessment”). The Risk Assessment shall—

(A) assess the potential implications of the covered artificial intelligence system on freedom of expression, equal protection, privacy, and due process;

(B) account for the environment in which the covered artificial intelligence system will be deployed, including its interactions with other artificial intelligence tools, programs, and systems that collect personally identifiable information; and

(C) include steps to mitigate and track any risks identified in the assessment.

(2) ARTIFICIAL INTELLIGENCE SYSTEM IMPACT ASSESSMENT.—Each covered agency shall conduct an Artificial Intelligence System Impact Assessment (“Impact Assessment”), no less than once per year, to assess the degree to which a covered artificial intelligence system remains compliant with the constraints and metrics established in the Risk Assessment. The Impact Assessment shall be based on outcomes, impacts, and metrics collected during system use, and shall determine if the existing validation processes should be improved.

(d) NOTICE OF DISCONTINUATION.—Within one year of discontinuing use of any non-public or classified covered artificial intelligence system, a covered agency shall

consider providing notice to the public that the covered artificial intelligence system has been discontinued.

(e) REPORT TO CONGRESS.—The head of each covered agency shall, within 90 days of the date of this Act, submit to Congress a report identifying any additional resources, including staff, needed to carry out the requirements of this section.

This section should be cross-referenced with the recommendation to create a task force to assess the privacy and civil rights and civil liberties implications of AI and emerging technologies, as the definition of a “covered artificial intelligence system” relies on the work of the task force.

Recommendation Set 2: Develop & Test Systems per Goals of Privacy Preservation and Fairness.

Establish third-party testing center(s) to allow independent, third-party testing of national security-related AI systems that could impact U.S. persons.

Require the Department of Justice (DOJ), in consultation with the Privacy and Civil Liberties Oversight Board (PCLOB), to develop binding guidance for the use of third-party testing (e.g., thresholds for high-consequence systems or unprecedented factors) of AI systems.

SEC. ____.—THIRD PARTY TESTING OF ARTIFICIAL INTELLIGENCE SYSTEMS.—

(a) IN GENERAL.—Not later than one year after the date of enactment of this Act, the Director of the National Institute of Standards and Technology shall establish an accreditation program for Third Party Independent Artificial Intelligence Testing Laboratories, as set forth in this section, to conduct independent testing of artificial intelligence systems for covered agencies to assess potential privacy, civil rights, and civil liberties impacts of such systems on U.S. persons.

(b) ARTIFICIAL INTELLIGENCE SYSTEMS REQUIRING TESTING.—The Privacy and Civil Liberties Oversight Board and the Department of Justice shall, in consultation with Privacy and Civil Liberties officers of the covered agencies, propose criteria for when an artificial intelligence system warrants third-party testing for privacy, civil liberties, and civil rights implications for U.S. Persons. Covered agencies shall adopt this criteria, as described in subsection (e).

(c) COVERED AGENCIES.—For the purposes of this section, covered agencies are the elements of the Intelligence Community (as defined in section 3003(4) of title 50, United States Code, and coordinated by the Office of the Director of National Intelligence), the Department of Homeland Security, and the Federal Bureau of Investigation.

(d) ACCREDITATION OF THIRD PARTY ARTIFICIAL INTELLIGENCE TESTING LABORATORIES.—Accreditation of Third Party Artificial Intelligence Testing Laboratories shall be done through the National Institute of Standards and Technology’s National Voluntary Laboratory Accreditation Program (“NVLAP”). In accordance with current NVLAP processes, the National Institute of Standards and Technology shall determine and maintain the authoritative list for approved Third Party Artificial Intelligence Testing Laboratories.

(e) INDEPENDENT TESTING REQUIRED.—Upon the approval of Third Party Artificial Intelligence Testing Laboratories as outlined in subsection (d), a covered agency, prior to procuring or fielding an artificial intelligence system requiring testing, shall institute independent third party testing of the system to assess performance of the system according to attributes listed in section 22A of the National Institute of Standards and Technology Act.

(f) SCOPE OF TESTING.—Each independent Third Party Artificial Intelligence Testing Laboratory accredited pursuant to subsection (d) shall—

(1) utilize metrics relevant to the mission and authorities of the agency that intends to field the artificial intelligence system;

(2) develop approaches to test—

(A) the software product, as installed in a test facility; and

(B) relevant cloud-based services.

(3) establish binding data agreements that enable the agency and other stakeholders to share confidential and proprietary data with the testing entity without fear of inappropriate disclosure; and

(4) collaborate with the covered agency that is seeking testing to reach consensus on appropriate protocols and approaches for handling test data, test results, and analyses.

Recommendation Set 4: Strengthen Oversight and Governance Mechanisms to Address Current and Evolving Concerns.

Strengthen the Privacy and Civil Liberties Oversight Board’s (PCLOB) ability to provide meaningful oversight and advice to the federal government’s use of AI-enabled technologies for counterterrorism purposes.

SEC.____.—OVERSIGHT OF FEDERAL GOVERNMENT USE OF ARTIFICIAL INTELLIGENCE-ENABLED SYSTEMS FOR COUNTERTERRORISM PURPOSES.—

(a) AMENDMENTS TO AUTHORITIES AND RESPONSIBILITIES OF THE PRIVACY AND CIVIL LIBERTIES OVERSIGHT BOARD.—Section 2000ee of title 42, United States Code, is amended—

(1) in paragraph (2) of subsection (d), by—

(A) striking “and” at the end of subparagraph (B);

(B) redesignating subparagraph (C) as subparagraph (D); and

(C) adding a new subparagraph (C), as follows:

“(C) the development and use of artificial intelligence-enabled technologies for counterterrorism purposes; and”;

(2) in subparagraph (1)(A) of subsection (g), by striking the semicolon and adding the following: “and information about artificial intelligence-enabled technologies proposed to be acquired or fielded in the Federal Government (such as documentation of data collection, disclosure and consent processes for artificial intelligence-enabled tools and programs, documentation of models used and supporting training and testing, and any repurposing);”

(b) AMENDMENTS TO AUTHORITIES AND RESPONSIBILITIES OF PRIVACY AND CIVIL LIBERTIES OFFICERS.—Section 2000ee-1 of title 42, United States Code, is amended—

(1) in subsection (a), by—

(A) redesignating paragraphs (3) and (4) as paragraphs (4) and (5); and

(B) inserting a new paragraph (3), as follows:

“(3) provide prior notice to the Privacy and Civil Liberties Oversight Board of the fielding or repurposing of an artificial intelligence-enabled system (including a classified system) that could have an impact on privacy, civil liberties, or civil rights, and provide access to associated impact statements, including System of Record Notices, Privacy Impact Assessments, and Civil Rights and Civil Liberties Impact Assessments;” and

(2) in subsection (d), by striking the semicolon in paragraph (1) and inserting the following: “(including information described in paragraph (a)(3));”.

(c) SELF-ASSESSMENT BY PRIVACY AND CIVIL LIBERTIES OVERSIGHT BOARD.—Not later than 270 days after the date of the enactment of this act, the Privacy and Civil Liberties Oversight Board shall conduct and provide to Congress a self-assessment of any change in resources and organizational structure that may be required to carry out the artificial intelligence-related mission required by this section.

Empower DHS Offices of Privacy and Civil Rights and Civil Liberties.

SEC. ____.—ENHANCED OVERSIGHT OF ARTIFICIAL INTELLIGENCE-ENABLED SYSTEMS AT THE DEPARTMENT OF HOMELAND SECURITY.—

(a) AMENDMENT TO DUTIES AND RESPONSIBILITIES OF CIVIL RIGHTS AND CIVIL LIBERTIES OFFICER.—Section 345 of title 6, United States Code, is amended in paragraph (a)(5), by—

- (1) striking the final “and” in subparagraph (A);
- (2) redesignating subparagraph (B) as subparagraph (C); and
- (3) adding a new subparagraph (B), as follows:

“(B) ensure that the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, provide appropriate consideration to the privacy, civil rights, and civil liberties impacts of such systems; and”.

(b) AMENDMENT TO DUTIES AND RESPONSIBILITIES OF CHIEF PRIVACY OFFICER.—Section 142 of title 6, United States Code, is amended in paragraph (a)(5), by—

- (1) striking the final “and” in subparagraph (A);
- (2) redesignating subparagraph (B) as subparagraph (C); and
- (3) adding a new subparagraph (B), as follows:

“(B) ensure that the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, provide appropriate consideration to the privacy, civil rights, and civil liberties impacts of such systems; and”.

(c) ENHANCED PROCEDURES FOR CONSIDERATION OF PRIVACY AND CIVIL LIBERTIES ISSUES.—Not later than 270 days after the date of the enactment of this Act—

(1) the Secretary of Homeland Security shall revise the legal and approval processes for the procurement and use of artificial intelligence-enabled systems, including associated data of machine learning systems, to ensure that full consideration is given, with the participation of the Department's Chief Privacy Officer and the Officer for Civil Rights and Civil Liberties, to the privacy, civil rights, and civil liberties impacts of such systems; and

(2) the Department's Chief Privacy Officer and the Officer for Civil Rights and Civil Liberties shall report to Congress on any additional staffing or funding resources that may be required to carry out the requirements of this section.

Establish a task force to assess the privacy and civil rights and civil liberties implications of AI and emerging technologies.

SEC. ____.—TASK FORCE ON ORGANIZATIONAL STRUCTURE FOR ARTIFICIAL INTELLIGENCE GOVERNANCE AND OVERSIGHT.—

(a) ESTABLISHMENT.—Not later than 90 days after the date of the enactment of this Act, the President shall appoint a task force to assess the privacy, civil rights, and civil liberties implications of artificial intelligence and emerging technologies. This includes identifying policy and legal gaps and making recommendations to ensure that uses of artificial intelligence and associated data in U.S. government operations comport with freedom of expression, equal protection, privacy, and due process. The task force shall—

(1) assess existing policy and legal gaps for current AI applications and emerging technologies, and make recommendations for—

(A) legislative and regulatory reforms on the development and fielding of AI and emerging technologies; and

(B) institutional changes to ensure sustained assessment and recurring guidance on privacy and civil liberties implications of AI applications and emerging technologies.

(b) MEMBERSHIP OF TASK FORCE.—

(1) The task force shall include—

(A) the Attorney General or his or her designee;

(B) the Director of the Office of Management and Budget or his or her designee;

(C) the Director of the National Institute of Standards and Technology or his or her designee;

(D) the Comptroller General or his or her designee;

(E) the Inspectors General for the following agencies:

(i) the Department of State;

(ii) the Department of the Treasury;

(iii) the Department of Defense;

(iv) the Department of Justice;

(v) the Department of Health and Human Services;

(vii) the Department of Homeland Security;

(viii) the Office of the Director of National Intelligence; and

(ix) the Central Intelligence Agency.

(F) the chief privacy and civil liberties officers of each agency described in subparagraph (E);

(G) the Chair of the Privacy and Civil Liberties Oversight Board;

(H) the Chair of the National Artificial Intelligence Advisory Committee's Subcommittee on Artificial Intelligence and Law Enforcement; and

(I) representatives from civil society, including organizational leaders with expertise in technology, privacy, civil liberties, and civil rights, representatives from industry, and representatives from academia, as appointed by the President.

(2) TASK FORCE CHAIR AND VICE CHAIR.—The President shall designate a Chair and Vice Chair of the task force from among its members.

(c) RESPONSIBILITIES OF TASK FORCE.—The task force established pursuant to subsection (a) shall—

(1) conduct an assessment and make recommendations to Congress and to the President to ensure that the development and fielding of artificial intelligence and other emerging technologies by the Federal Government provides protections for the privacy, civil liberties, and civil rights of U.S. persons as appropriately balanced against critical law enforcement and national security needs;

(2) issue criteria for identifying qualified artificial intelligence systems and significant system refreshes requiring Artificial Intelligence Risk Assessment Reports and Artificial Intelligence Impact Assessments, under section [XX] of this Act;

(3) recommend baseline standards for Federal Government use of biometric identification technologies, including, but not limited to, facial recognition, voiceprint, gait recognition, and keyboard entry technologies;

(4) recommend proposals to address any gaps in Federal law or regulation with respect to facial recognition technologies in order to enhance protections of privacy, civil liberties, and civil rights of U.S. persons;

(5) recommend best practices and contractual requirements to strengthen protections for privacy, information security, fairness, non-discrimination, auditability, and accountability in artificial intelligence systems and technologies and associated data procured by the federal government;

(6) consider updates to and reforms of government data privacy and retention requirements to address implications to privacy, civil liberties, and civil rights;

(7) assess ongoing efforts to regulate commercial development and fielding of artificial intelligence and associated data in light of privacy, civil liberties, and civil rights implications, and as appropriate, consider and recommend institutional or organizational changes to facilitate applicable regulation; and

(8) assess the utility of establishing a new organization within the Federal Government to provide ongoing governance for and oversight over the fielding of artificial intelligence technologies by Federal agencies as technological capabilities evolve over time.

(d) ORGANIZATIONAL CONSIDERATIONS.—In conducting the assessment required by subsection (c)(7), the task force shall consider—

(1) the organizational placement, structure, composition, authorities, and resources that a new organization would require to provide ongoing guidance and baseline standards for—

(A) the Federal Government’s development, acquisition, and fielding of artificial intelligence systems to ensure they comport with privacy, civil liberties, and civil rights and civil liberties law, to include guardrails for their use and to disallow outcomes to be incorporated in policy and embedded in system development; and

(B) providing transparency to oversight entities and the public regarding the Federal Government’s use of artificial systems and the performance of those systems.

(2) the existing interagency and intra-agency efforts to address AI oversight;

(3) the need for and scope of national security carve outs, and any limitations or protections that should be built into any such carve outs; and

(4) the research, development, and application of new technologies to mitigate privacy and civil liberties risks inherent in artificial intelligence systems.

(e) REPORTING.—

(1) Not later than 180 days of establishment, the task force shall issue a report to Congress and the President with its legislative and regulatory recommendations. The task force shall provide periodic updates to the President and the Congress.

(2) Within a year of its establishment, the task force shall issue a report to the President and the Congress with its assessment on organizational considerations, to include any recommendations for organizational changes.

CHAPTER 10: THE TALENT COMPETITION

Blueprint for Action

Recommendation: Pass a National Security Immigration Act.

1) Grant Green Cards to All Students Graduating with STEM PhDs from Accredited American Universities.

2) Double the Number of Employment Based Green Cards.

3) Create an Entrepreneur Visa.

4) Create an Emerging and Disruptive Technology Visa.

NATIONAL SECURITY IMMIGRATION ACT OF 2021

SECTION. 1.—SHORT TITLE.—This Act may be cited as the “National Security Immigration Act of 2021.”

SEC. 2.—GREEN CARDS FOR STUDENTS GRADUATING FROM ACCREDITED AMERICAN UNIVERSITIES WITH DOCTORATES IN THE FIELDS OF SCIENCE, TECHNOLOGY, ENGINEERING, AND MATHEMATICS.—Section 1151 of title 8, United States Code, is amended in subsection (b)(1), by adding a new subparagraph (F), as follows:

“(F) Aliens who have been awarded doctoral degrees in the fields of science, technology, engineering, and mathematics by accredited universities in the United States.”

SEC. 3.—INCREASED AUTHORIZATION FOR EMPLOYMENT-BASED IMMIGRATION.—Section 1151 of title 8, United States Code, as amended by section 2, is further amended in subsection (d)(1)(A) by striking “140,000” and inserting “280,000”.

SEC. 4.—ENTREPRENEUR VISAS FOR HIGH PRIORITY SCIENCE AND TECHNOLOGY FIELDS AS DETERMINED BY NATIONAL SCIENCE FOUNDATION.—Section 1153 of title 8, United States Code, is amended in subsection (b)(5)—

(1) By redesignating subparagraphs (C) and (D) as subparagraphs (D) and (E); and

(2) By adding a new subparagraph (C), as follows:

“(C) PRIORITY FOR ENTREPRENEURS IN CERTAIN SCIENCE AND TECHNOLOGY FIELDS.—

“(i) Priority under this section shall be given to qualified immigrants who engage in new commercial enterprises in high priority science and technology fields, including artificial intelligence-enabled technology fields, as determined by the National Science Foundation.

“(ii) A qualified immigrant under this paragraph section shall not be required to meet the capital investment requirement in clause (A)(i) if the qualified immigrant is one of the principal organizers and operators of a new commercial enterprise described in clause (i).”

SEC. 5.—VISA FOR EMERGING AND DISRUPTIVE TECHNOLOGIES.—Section 1151 of title 8, United States Code, as amended by Sections 2 and 3, is further amended in subsection (b)(1), by adding a new clause (G), as follows:

“(G) Aliens who are students, researchers, entrepreneurs, and technologists in critical emerging and disruptive technology fields, as determined by the National Science Foundation.”

SEC. 6.—DETERMINATIONS BY THE NATIONAL SCIENCE FOUNDATION.—Not later than 180 days after the date of the enactment of this Act, and every three years thereafter, the National Science Foundation shall publish a list of—

(1) high priority science and technology fields in which qualified immigrants will be eligible for consideration for entrepreneur visas under section 1153(b)(5)(C) of title 8, United States Code, as amended; and

(2) critical emerging and disruptive technology fields in which qualified immigrants will be eligible for consideration for student, researcher, and entrepreneur visas under section 1151(b)(1)(G) of title 8, United States Code, as amended.

CHAPTER 11: ACCELERATING AI INNOVATION

Blueprint for Action

Recommendation: Scale and Coordinate Federal AI R&D Funding.

Component 1: Establish a National Technology Foundation.

THE NATIONAL TECHNOLOGY FOUNDATION ACT OF 2021

SECTION 1.—SHORT TITLE.—This Act may be cited as the “National Technology Foundation Act of 2021.”

SEC. 2.—ESTABLISHMENT OF NATIONAL TECHNOLOGY FOUNDATION.—There is established in the executive branch of the Government an independent agency to be known as the National Technology Foundation (hereinafter referred to as the “Foundation”). The Foundation shall consist of a National Technology Board (hereinafter referred to as the “Board”) and a Director of the Foundation (hereinafter referred to as the “Director”).

SEC. 3.—NATIONAL TECHNOLOGY BOARD.—

(a) The Board shall consist of twenty-four members to be appointed by the President and of the Director ex officio. In addition to any powers and functions otherwise granted to it by this chapter, the Board shall establish the policies of the Foundation, within the framework of applicable national policies as set forth by the President and the Congress.

(b) The term of office of each member of the Board shall be six years; except that any member appointed to fill a vacancy occurring prior to the expiration of the term for which his predecessor was appointed shall be appointed for the remainder of such term. Any person, other than the Director, who has been a member of the Board for twelve consecutive years shall thereafter be ineligible for appointment during the two-year period following the expiration of such twelfth year.

SEC. 4.—DIRECTOR OF THE FOUNDATION.—The Director shall be appointed by the President, by and with the advice and consent of the Senate. Before any person is appointed as Director, the President shall afford the Board an opportunity to make recommendations to the President with respect to such appointment. The Director shall receive basic pay at the rate provided for level II of the Executive Schedule under Section 5313 of title 5,

United States Code, and shall serve for a term of six years unless sooner removed by the President.

SEC. 5.—DEPUTY DIRECTOR OF THE FOUNDATION.—The Deputy Director (hereinafter referred to as the “Deputy Director”) shall be appointed by the President, by and with the advice and consent of the Senate. Before any person is appointed as a Deputy Director, the President shall afford the Board and the Director an opportunity to make recommendations to the President with respect to such appointment. The Deputy Director shall receive basic pay at the rate provided for level III of the Executive Schedule under section 5314 of title 5, United States Code, and shall perform such duties and exercise such powers as the Director may prescribe. The Deputy Director shall act for, and exercise the powers of, the Director during the absence or disability of the Director, or in the event of a vacancy in the office of Director.

SEC. 6.—GENERAL AUTHORITY OF THE FOUNDATION.—

(a) The Foundation shall have the authority, within the limits of available appropriations, to do all things necessary to carry out the provisions of this chapter, including, but without being limited thereto, to—

(1) distribute other payments for research and development in priority technology areas through grants, cooperative agreements, and contracts awarded to academic and private sector researchers, nonprofits, and consortia through competitive processes without regard to the provisions of sections 3324(a) and (b) of title 31, United States Code;

(2) establish an innovation unit in which independent program managers, brought into the Foundation on the basis of term appointments, fund proposals from both industry and academia to advance solutions to forward-looking research questions in priority technology areas;

(3) organize prize competitions to catalyze research around significant technology challenge problems;

(4) manage national technology resources, infrastructure, and initiatives that are assigned to the Foundation by statute or executive order;

(5) promote the commercialization of new technologies in priority technology areas and the transfer of such technologies to Federal, State and local government entities; and

(6) serve as a focal point for international research and development collaboration and standards-setting dialogues in priority technology areas.

SEC. 7.—PRIORITY TECHNOLOGY AREAS.—

(a) CORE DIRECTORATES.—The Foundation shall be organized into a set of core directorates, each dedicated to advancing fundamental research into a priority technology area.

(b) PRIORITY TECHNOLOGY AREAS.—Priority technology areas shall include—

- (1) artificial intelligence;
- (2) biotechnology;
- (3) quantum computing;
- (4) semiconductors and advanced hardware;
- (5) robotics and autonomy;
- (6) fifth-generation and advanced networking;
- (7) advanced manufacturing;
- (8) energy technology; and
- (9) any other technology area designated by the Congress or the Board.

(c) REVIEW OF KEY TECHNOLOGY FOCUS AREAS AND SUBSEQUENT LISTS.—

(1) ADDING OR DELETING KEY TECHNOLOGY FOCUS AREAS.—Beginning on the date that is four years after the date of enactment of this Act and every four years thereafter, the Director, acting through the Deputy Director shall—

(A) review the list of key technology focus areas, in consultation with the Board; and

(B) as part of that review, may add or delete key technology focus areas if the competitive threats to the United States have shifted and whether the United States or other nations have advanced or fallen behind in a technological area.

(2) LIMIT ON KEY TECHNOLOGY FOCUS AREAS.—Not more than ten key technology focus areas shall be included on the list of key technology focus areas at any time.

(3) UPDATING FOCUS AREAS AND DISTRIBUTION.—Upon the completion of each review under this subsection, the Director shall make the list of key technology focus areas readily available and publish the list in the Federal Register, even if no changes have been made to the prior list.

SEC. 8.—ADMINISTRATIVE MATTERS.—

(a) HIRING AUTHORITY.—

(1) PRIORITY TECHNOLOGY EXPERTS.—The Director shall have the authority to carry out a program of personnel management authority for the Foundation in the same manner, and subject to the same requirements, as the program of personnel management authority authorized for the Director of the Defense Advanced Research Projects Agency under section 1599h(a)(2) of title 10, United States Code, for the Defense Advanced Research Projects Agency.

(2) HIGHLY QUALIFIED EXPERTS.—In addition to the authority provided under subsection (A), the Director shall have the authority to carry out a program of personnel management authority for the Foundation in the same manner, and subject to the same requirements, as the program to attract highly qualified experts carried out by the Secretary of Defense under section 9903 of title 5, United States Code.

(3) ADDITIONAL HIRING AUTHORITY.—To the extent needed to carry out the duties of the Foundation, the Director shall utilize hiring authorities under section 3372 of title 5, United States Code, to staff the Foundation with employees from other Federal agencies, State and local governments, Indian tribes and tribal organizations, institutions of higher education, and other organizations, as described in that section, in the same manner and subject to the same conditions.

(b) EMPLOYMENT AND COMPENSATION OF CERTAIN PERSONNEL.—

(1) PROGRAM MANAGERS.—The employees of the Foundation may include program managers, who shall perform a role similar to program managers employed by the Defense Advanced Research Projects Agency, for the oversight and selection of programs supported by the Foundation.

(2) COMPENSATION OF MEMBERS OF BOARD.—The members of the Board shall be entitled to receive compensation for each day engaged in the business of the Foundation at a rate fixed by the Chairman but not exceeding the maximum rate payable under section 5376 of title 5, United States Code, and shall be allowed travel expenses as authorized by 5703 of title 5, United States Code. For the purposes of determining the payment of compensation under this subsection, the time spent in travel by any member of the Board shall be deemed as time engaged in the business of the Foundation. Members of the Board and

members of special commissions may waive compensation and reimbursement for traveling expenses.

SEC. 9.—INTERNATIONAL COOPERATION.—

(a) INTERNATIONAL AUTHORITY.—The Foundation is authorized to cooperate in any international technology activities consistent with the purposes of this Act and to expend for such international technology activities such sums within the limit of appropriated funds as the Foundation may deem appropriate.

(b) CONTRACTS AND ARRANGEMENTS.—

(1) The authority to enter into contracts or other arrangements with organizations or individuals in foreign countries and with agencies of foreign countries, as provided in section 1870(c) of title 42, United States Code, and the authority to cooperate in international scientific or engineering activities as provided in subsection (a) of this section, shall be exercised only with the approval of the Secretary of State, to the end that such authority shall be exercised in such manner as is consistent with the foreign policy objectives of the United States.

(2) If, in the exercise of the authority referred to in paragraph (1) of this subsection, negotiation with foreign countries or agencies thereof becomes necessary, such negotiation shall be carried on by the Secretary of State in consultation with the Director.

SEC. 10.—SECURITY PROVISIONS.—

(a) RESEARCH RELATED TO NUCLEAR ENERGY.— The Foundation shall not support any research or development activity in the field of nuclear energy, nor shall it exercise any authority pursuant to section 1870(e) of title 42, United States Code, in respect to that field, without first having obtained the concurrence of the Secretary of Energy that such activity will not adversely affect the common defense and security. To the extent that such activity involves restricted data as defined in the Atomic Energy Act of 1954, the provisions of that Act regarding the control of the dissemination of restricted data and the security clearance of those individuals to be given access to restricted data shall be applicable. Nothing in this chapter shall supersede or modify any provision of the Atomic Energy Act of 1954.

(b) RESEARCH RELATION TO NATIONAL DEFENSE.—

(1) In the case of priority technology area research activities under this Act in connection with matters relating to the national defense, the Secretary of Defense shall establish such security requirements and safeguards, including restrictions with respect to access to information and property, as the Secretary of Defense deems necessary.

(2) Any agency of the Government exercising investigatory functions otherwise within its jurisdiction is authorized to make such investigations and reports as may be requested by the Foundation in connection with the enforcement of security requirements and safeguards, including restrictions with respect to access to information and property, established under paragraph (1) of this subsection.

SEC. 11.—REPORTS.—

(a) INITIAL REPORT.—Not later than one year after the date of enactment of this Act, the Director shall transmit to the Committee on Commerce, Science, and Transportation of the Senate and the Committee on Science, Space, and Technology of the House of Representatives a report regarding the establishment of the Foundation. The report shall include an assessment of the priority technology focus areas as defined in this Act and of authorities that conflict with the National Science Foundation.

(b) ANNUAL REPORTS.—

(1) The Board shall submit to the President and the Congress no later than January 15 of each even numbered year, a report on indicators of the state of the priority technology areas in the United States, as defined in this Act.

(2) The Board shall render to the President and the Congress reports on specific, individual policy matters within the authority of the Foundation (or otherwise as requested by the Congress or the President) related to priority technology areas, as the Board, the President, or the Congress determines the need for such reports.

SEC. 12.—AUTHORIZATION OF APPROPRIATIONS.—

(a) INITIAL APPROPRIATION.—To enable the Foundation to carry out its powers and duties, including the establishment of a physical location, there is authorized to be appropriated to the Foundation \$30,000,000 for the first fiscal year following the enactment of this Act. Appropriations made pursuant to the authority provided in this subsection shall remain available for obligation, for expenditure, or for obligation and expenditure until expended for the Foundation's initial administrative costs and salaries and expenses.

(b) ANNUAL APPROPRIATION.—There are authorized to be appropriated for the Foundation, in addition to the appropriation provided in subsection (a) of this section and any other funds made available to the Foundation, a total of \$51,000,000,000 for fiscal years 2022 through 2026, of which—

(A) \$1,000,000,000 is authorized for fiscal year 2022;

(B) \$5,000,000,000 is authorized for fiscal year 2023;

(C) \$10,000,000,000 is authorized for fiscal year 2024;

(D) \$15,000,000,000 is authorized for fiscal year 2025; and

(E) \$20,000,000,000 is authorized for fiscal year 2026.

The Commission acknowledges additional authorities may be required to establish the NTF, including administrative, financial, and educational authorities mirroring those of the National Science Foundation, and that amendments to the NSF's statutory authorities may be required to alleviate duplication of duties. The Commission is ready to work with Congress to address such provisions.

Component 4: Invest in Talent that Will Transform the Field.

Direct and fund establishment of an AI Innovator Award.

Direct and fund establishment of a team-based AI research award.

SEC. ____.—ARTIFICIAL INTELLIGENCE AWARD PROGRAM.—

(a) ARTIFICIAL INTELLIGENCE INNOVATOR AWARD.—

(1) IN GENERAL.—The Director of the National Science Foundation shall partner with a nonprofit organization as described in subsection (c) to establish an Artificial Intelligence Innovator Award program to recognize and support the research of leaders in the field of artificial intelligence.

(2) ARTIFICIAL INTELLIGENCE INNOVATOR AWARD RECIPIENTS.—The Artificial Intelligence Award Selection Committee as described in subsection (d) shall select no fewer than 10 and no more than 20 award recipients each year. Recipients shall be selected for five-year, renewable award terms, based on a proven track record of prior innovation, a proposed general research program, a commitment to spend 75 percent of the recipients' time on research, and the committee's assessment of the potential of the research to generate breakthroughs in the area of artificial intelligence. Award amounts shall be determined by the selection committee with the objective of covering the full salary and benefits of the researcher and the cost of associated support staff and research equipment.

(b) ARTIFICIAL INTELLIGENCE TEAM AWARD.—

(1) IN GENERAL.—The Director of the National Science Foundation shall partner with a nonprofit organization as described in subsection (c) to establish an Artificial Intelligence Team Award program to support interdisciplinary research directed at applying artificial intelligence to solve complex problems or pursuing use-inspired basic research efforts to advance a fundamental understanding of the science of artificial intelligence in a manner that provides a significant benefit to society.

(2) ARTIFICIAL INTELLIGENCE TEAM AWARD RECIPIENTS.—The Artificial Intelligence Innovator Awards Selection Committee as described in paragraph (d) shall select no fewer than five and no more than 10 team recipients each year. Recipients shall be selected for five-year, nonrenewable terms, based on team qualifications, commitment to multi-disciplinary approaches, and innovative research proposals. Award amounts shall be determined by the selection committee with the objective of covering the cost of carrying out the proposed research proposal.

(c) NONPROFIT ORGANIZATION PARTNER.—The National Science Foundation shall partner with a nonprofit organization active in the field of computer science and artificial intelligence that maintains the requisite expertise and connections to the artificial intelligence research community to identify promising talent and invest in innovative ideas and to manage the award programs described in subsections (a) and (b), including to administer the programs and arrange the annual meeting.

(d) ARTIFICIAL INTELLIGENCE AWARD SELECTION COMMITTEE.—Recipients of the Artificial Intelligence Innovator Award and the Artificial Intelligence Team Award shall be selected by a rotating committee of artificial intelligence experts known as the Artificial Intelligence Award Selection Committee. The Committee shall consist of members chosen for their first-hand experience in artificial intelligence research and their familiarity with the frontiers of the field. Committee member selection shall be made by the nonprofit organization partner identified under subsection (c), in consultation with the Director of the National Science Foundation or designee.

(e) ANNUAL MEETING.—The Director of the National Science Foundation shall sponsor an annual meeting of recipients of the Artificial Intelligence Innovator Award and the Artificial Intelligence Team Award, at which the award recipients shall share information on the progress of their work.

(f) OTHER SOURCES OF FUNDING.—Nothing in this section shall be interpreted to preclude a recipient of an Artificial Intelligence Innovator Award or an Artificial Intelligence Team Award from pursuing supplemental government research grant or other research support provided by individuals, nonprofits and corporations, provided that such additional funding does not interfere with the recipient's commitment to the research program or require the assignment of ownership of intellectual property in a manner that would be inconsistent with the provisions of the Bayh-Dole Act, Public Law 96-517.

(g) INDEPENDENT REVIEW.—The Director of the National Science Foundation shall engage an independent entity to conduct a review to assess the successes and failures of the awards program authorized by this section, evaluate the impact of the funding level and award term on the research conducted by participants, and recommend any needed changes to the program (including any expansion or contraction in the number of

awards). The findings of the independent review shall be delivered to Congress not later than seven years after the commencement of the program.

(h) AUTHORIZATION OF APPROPRIATION.—

(1) There is authorized to be appropriated for each of the fiscal years 2022 through 2028 \$125,000,000 for the Artificial Intelligence Innovator Award.

(2) There is authorized to be appropriated for the Artificial Intelligence Team Award—

(A) \$50,000,000 for fiscal year 2022;

(B) \$100,000,000 for fiscal year 2023;

(C) \$150,000,000 for fiscal year 2024;

(D) \$200,000,000 for fiscal year 2025; and

(E) \$250,000,000 for fiscal years 2026 through 2028.

Recommendation: Leverage Both Sides of the Public-Private Partnership.

Component 2: Form a Network of Regional Innovation Clusters Focused on Strategic Emerging Technologies.

SEC. ____.—ESTABLISHMENT OF A NATIONAL NETWORK FOR REGIONAL INNOVATION IN EMERGING TECHNOLOGIES.—

(a) ESTABLISHMENT OF NATIONAL PROGRAM OFFICE.—The Secretary of Commerce shall establish, within the National Institute of Standards and Technology, a National Program Office for Regional Innovation in Emerging Technologies (referred to in this section as the ‘National Program Office’).

(b) DUTIES AND RESPONSIBILITIES.—The National Program Office, in coordination with representatives of Federal agencies with experience in and missions related to emerging technologies, shall—

(1) oversee the planning, development, management, and coordination of a National Network for Regional Innovation in Emerging Technologies (referred to in this section as the “National Network”);

(2) develop, not later than one year after the date of enactment, and update not less frequently than once every three years thereafter, a strategic plan to guide the development of the National Network to include identification of priority emerging technologies critical to national security or national competitiveness;

(3) use a competitive process to designate and provide financial assistance to regional innovation clusters that enable United States leadership in emerging technologies and support regional economic development throughout the United States;

(4) establish within each regional innovation cluster in the National Network a Technology Research Center for the purpose of facilitating collaboration among regional innovation cluster participants;

(5) establish such procedures, processes, and criteria as may be necessary and appropriate to coordinate the activities of the National Network and to maximize participation in and coordination with the National Network by Federal agencies that field or operate systems that incorporate emerging technologies;

(6) establish a clearinghouse of public information related to the activities of the National Network; and

(7) act as a convener of the National Network.

(c) DESIGNATION OF AND FINANCIAL ASSISTANCE IN SUPPORT OF REGIONAL INNOVATION CLUSTERS.—The National Program Office shall use a competitive process to designate and provide financial assistance to regional innovation clusters based on the following criteria:

(1) the equitable distribution of regional innovation clusters throughout the United States, taking into account factors such as proximity to the research and development facilities of Federal agencies, the level of support from state and local governments, the presence of and value proposition for leading firms and research institutions in relevant fields, and the size and education level of the local workforce;

(2) the capacity of regional innovation clusters to support the research, development, and commercialization of specific emerging technologies in areas that are critical to United States national competitiveness; and

(3) the clear potential for future development of regional innovation clusters that are not yet established technology hubs.

(d) TECHNOLOGY RESEARCH CENTERS.—The National Program Office shall establish within each regional innovation cluster in the National Network a Technology Research Center for the purpose of facilitating collaboration between regional innovation cluster participants. The Technology Research Centers shall—

(1) form sustained partnerships with anchor institutions in the region;

(2) host researchers on temporary assignments from Federal agencies, establish talent exchanges with local firms and research institutions, and fund multi-year, post-doctoral fellowships for the commercialization of research;

(3) host program managers from Federal agencies responsible for transitioning basic research into commercially viable technologies, identifying national security use cases and end users within the Federal Government, and initiating new Federal Government contracts to support technology transition;

(4) facilitate low cost access by regional innovation cluster participants to computing resources, curated datasets, testing infrastructure and ranges, and other research and development facilities owned or operated by the Federal government;

(5) establish intellectual property sharing agreements with regional innovation cluster participants to encourage Federal government adoption of commercial technologies; and

(6) when appropriate, provide for the publication of research in the open-source domain to encourage advances in the science and technology community more broadly.

(e) OTHER MATTERS.—

(1) RECOMMENDATIONS.—In developing and updating the strategic plan under subsection (b)(2), the National Program Office shall solicit recommendations and advice from a wide range of stakeholders, including industry, small and medium-sized enterprises, research universities, community colleges, state and local elected officials, and other relevant organizations and institutions on an ongoing basis.

(2) REPORT TO CONGRESS.—Upon completion of the strategic plan required by subsection (b)(2) or an update thereof, the National Program Office shall transmit the strategic plan to the Committee on Commerce, Science, and Transportation of the Senate and the Committee on Science, Space, and Technology of the House of Representatives.

(3) DETAILEES.—Any Federal Government employee may be detailed to the National Program Office without reimbursement. Such detail shall be without interruption or loss of civil service status or privilege.

(f) DEFINITIONS.—

(1) REGIONAL INNOVATION CLUSTER.—The term “regional innovation cluster” means a geographically bounded network of similar, synergistic, or complementary entities that —

(A) are engaged in or with a particular industry sector and its related sectors;

(B) have active channels for business transactions and communication;

(C) share specialized infrastructure, labor markets, and services; and

(D) leverage the region’s unique competitive strengths to stimulate innovation and create jobs.

(2) EMERGING TECHNOLOGIES.—For the purposes of this section the term “emerging technologies” may include such technologies as artificial intelligence, microelectronics, quantum computing, biotechnology, any associated, enabling or successor technologies, or any technologies identified by the National Program Office to be critical to national security or national competitiveness.

(g) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated to the Secretary of Commerce to carry out this section \$5,000,000 for fiscal year 2022.

CHAPTER 14: TECHNOLOGY PROTECTION*Blueprint for Action*

Recommendation: Reform CFIUS for Emerging Technology Competition.

Amend CFIUS’ authorizing legislation to require competitors to disclose investments in “sensitive technologies” to CFIUS.

SEC. ____ . REVIEW OF SENSITIVE TRANSACTIONS INVOLVING COUNTRIES OF SPECIAL CONCERN.

(a) TECHNICAL AMENDMENTS.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by redesignating paragraphs (4), (5), (6), (7), (8), (9), (10), (11), (12), and (13) as paragraphs (5), (6), (7), (9), (10), (11), (12), (13), (15), and (16), respectively.

(b) DEFINITION OF COUNTRY OF SPECIAL CONCERN.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after paragraph (3) the following:

“(4) COUNTRY OF SPECIAL CONCERN.—The term “country of special concern” means any country that is—

“(A) subject to export restrictions pursuant to section 744.21 of title 15, Code of Federal Regulations;

“(B) determined by the Secretary of State to be a state sponsor of terrorism; or

“(C) determined by the Committee to have a demonstrated or declared strategic goal of acquiring a type of technology or infrastructure that would have an adverse impact on United States leadership in areas related to national security, and is specified in regulations prescribed by the Committee.”

(c) DEFINITION OF SENSITIVE TECHNOLOGY.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after redesignated paragraph (7) the following:

“(8) SENSITIVE TECHNOLOGY.—The term ‘sensitive technology’ means any technology that is determined by the Committee to be necessary for maintaining or increasing the technological advantage of the United States over countries of special concern with respect to national defense, intelligence, or other areas of national security, or gaining such an advantage over such countries with respect to national defense, intelligence, or other areas of national security in areas where such an advantage may not exist, and is not a critical technology as defined in paragraph (7) of this subsection, and is specified in regulations prescribed by the Committee.

(d) DEFINITION OF SENSITIVE TRANSACTION INVOLVING A COUNTRY OF SPECIAL CONCERN.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended by inserting after redesignated paragraph (13) the following:

“(14) SENSITIVE TRANSACTION INVOLVING A COUNTRY OF SPECIAL CONCERN.—The term ‘sensitive transaction involving a country of special concern’ means any investment in an unaffiliated United States business by a foreign person that—

“(A) is—

“(i) a national or a government of, or a foreign entity organized under the laws of, a country of special concern; or

“(ii) a foreign entity—

“(I) over which control is exercised or exercisable by a national or a government of, or by a foreign entity organized under the laws of, a country of special concern; or

“(II) in which the government of a country of special concern has a substantial interest; and

“(B) as a result of the transaction, could achieve—

“(i) influence, other than through voting of shares, on substantive decision making of the United States business regarding the use, development, acquisition, or release of sensitive technologies, as defined in this section; or—

“(ii) access to material nonpublic technical information related to sensitive technologies, as defined in this section, in the possession of the United States business.”

(e) DEFINITION OF COVERED TRANSACTIONS.—Section 721(a) of the Defense Production Act of 1950 (50 USC 4565(a)) is amended—

(1) in redesignated paragraph (5)(B)—

(A) in clause (iv)(I), by striking “or”;

(B) in clause (iv)(II), by striking the period and inserting “; or”; and

(C) by adding at the end the following:

“(III) a sensitive transaction involving a country of special concern.”

(2) by redesignating clause (v) as clause (vi) and inserting after clause (iv) the following:

“(v) Any sensitive transaction involving a country of special concern.”

(f) INFORMATION REQUIRED IN ANNUAL REPORT TO CONGRESS.—Section 721(m)(2) of the Defense Production Act of 1950 (50 USC 4565(m)(2)) is amended by adding at the end the following:

“(L) Identification of each country designated as a country of special concern along with an explanation of the rationale for such designation.

“(M) Identification of each technology designated as a sensitive technology along with an explanation of the rationale for such designation.”

(g) MANDATORY DECLARATIONS.—Section 721(b)(1)(C)(v)(IV)(bb)(AA) of the Defense Production Act of 1950 (50 USC 4565(b)(1)(C)(v)(IV)(bb)(AA)) is amended by inserting before the period “or is a sensitive transaction involving a country of special concern”.

(h) CONFORMING AMENDMENTS.—Title 50, United States Code, is amended—

(1) in section 4817(a)(1)(B) by striking “section 4565(a)(6)(A)” and inserting “section 4565(a)(7)(A)”;

(2) in section 4565(b)(4)(B)(ii) (section 721(b)(4)(B)(ii) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(ii)” and inserting “subsection (a)(5)(B)(ii)”;

(3) in section 4565(b)(1)(c)(v)(III)(bb)(AA) (section 721(b)(1)(c)(v)(III)(bb)(AA) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B) (iii)” and inserting “subsection (a)(5)(B)(iii)”;

(4) in section 4565(b)(1)(c)(v)(III)(bb)(BB) (section 721(b)(1)(c)(v)(III)(bb)(BB) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(iii)” and inserting “subsection (a)(5)(B)(iii)”;

(5) in section 4565(b)(1)(c)(v)(III)(cc) (section 721(b)(1)(c)(v)(III)(bb)(BB) of the Defense Production Act of 1950) by striking “subsection (a)(4)(B)(iii)(II)” and inserting “subsection (a)(5)(B)(iii)(II)”.

Recommendation: Build Capacity to Protect the Integrity of the U.S. Research Environment. Establish a government-sponsored independent entity focused on research integrity.

SEC. ____.—Establishment of University Affiliated Research Center Focused on Research Integrity.—

(a) AGREEMENT AUTHORIZED.—Not later than 180 days after the date of the

enactment of this Act, the Secretary of Defense, acting through the Under Secretary of Defense for Research and Engineering and in consultation with the Director of the Office of Science and Technology Policy and other appropriate members of the Federal research community, shall enter into an agreement with a college or university to establish a University Affiliated Research Center to act as a center of excellence on research integrity and provide information and advice on research security.

(b) RESEARCH PURPOSES.—The University Affiliated Research Center established pursuant to subsection (a) shall—

(1) Maintain open source materials to serve university vetting of international engagement and risk management, including databases and risk assessment tools;

(2) Provide tailored guidance to research organizations for decision support on matters related to research security and integrity;

(3) Conduct comprehensive studies and regular reports on the state of foreign influence on U.S. research;

(4) Undertake independent investigations on research integrity;

(5) Develop education materials and tools for U.S. universities to build annual training and compliance initiatives; and

(6) Manage dialogue with stakeholder communities and provide a venue for information sharing among research organizations and Federal agencies.

*Recommendation: Counter Foreign Talent Recruitment Programs.
Mandate and resource compliance operations.*

SEC. ____.—Enhanced Review of Risk Posed by Applicants for Federal Grants.—

(a) ENHANCED REVIEW REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Director of the Office of Management and Budget shall revise section 200.206 of Part 2 of the Code of Federal Regulations to ensure that Federal grant-making agencies maintain compliance operations to guard against malign foreign talent recruitment programs and to prescribe standardized disclosure and accountability measures to support such compliance operations.

(b) DEFINITION.—For the purposes of this section, a “malign foreign talent recruitment program” is an effort directly or indirectly organized, managed, or funded by a foreign government to recruit science and technology professionals or students (regardless of citizenship or national origin) engaged in research funded by a federal agency to share information with or otherwise act on behalf of such foreign government.

Amend the Foreign Agents Registration Act.

SEC. ____.—AMENDMENT TO FOREIGN AGENTS REGISTRATION ACT. —Section 611 of title 22, United States Code, is amended in paragraph (1) of subsection (c) by—

- (1) Striking “and” at the end of clause (iv); and
- (2) Inserting at the end a new clause (v), as follows:

“(v) directly or indirectly organizes, manages, or funds an effort to recruit science and technology professionals or students (regardless of citizenship or national origin) engaged in research funded by a Federal agency to share information with or otherwise act on behalf of a foreign government; and”.

CHAPTER 15: A FAVORABLE INTERNATIONAL TECHNOLOGY ORDER

Blueprint for Action

Recommendation: Develop and Implement a Comprehensive U.S. National Plan to Support International Technology Efforts.

Core Goal #1: Shape International Technical Standards.

Establish a grant program to enable small- and medium-sized U.S. AI companies to participate in international standardization efforts.

SEC. ____.—SUPPORT FOR INDUSTRY PARTICIPATION IN INTERNATIONAL STANDARDS ORGANIZATIONS.—

(a) IN GENERAL.—Not later than 180 days after the date of the enactment of this Act, the Administrator of the Small Business Administration shall establish a program to support participation by small business concerns in meetings and proceedings of international standards organizations in the development of voluntary technical standards.

(b) GRANTS AUTHORIZED.—In carrying out the program authorized by subsection (a), the Administrator shall award competitive, merit-reviewed grants, to small business concerns to cover the reasonable costs, up to a specified ceiling, of participation of employees of such businesses in meetings and proceedings of international standards organizations. Participation may include regularly attending meetings, contributing expertise and research, proposing new work items, volunteering for leadership roles such as convenors and editors, and being early adopters of emerging standards. Recipients of awards under this subsection shall not be required to provide a matching contribution.

(c) AWARD CRITERIA.—The Administrator may provide under this section a grant award to covered entities that:

- (1) demonstrate deep technical expertise in key emerging technologies, including Artificial Intelligence and related technologies;

(2) commit personnel with such expertise to regular participation in international bodies responsible for setting standards for such technologies over the period of the grant; and

(3) agree to participate in efforts to coordinate between the U.S. government and industry to ensure protection of national security interests in the setting of international standards.

(d) EVALUATION.—In issuing awards under this section, the Administrator shall coordinate with the Director of the National Institute of Standards and Technology who shall provide support in the assessment of technical expertise in emerging technologies and standards setting needs.

(e) DEFINITIONS.—In this section:

(1) ADMINISTRATOR.—The term “Administrator” means the Administrator of the Small Business Administration.

(2) COVERED ENTITY.—The term “covered entity” means a small business concern that is incorporated in and maintains a primary place of business in the United States.

(3) SMALL BUSINESS CONCERN.—The term “small business concern” has the same definition as set out in section 632 of title 15, United States Code.

(f) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated for fiscal year 2022 and each fiscal year thereafter \$1,000,000 to carry out the program authorized in this section.

Core Goal #2: Implement a Coordinated U.S. National Policy for the IDDI.

Create an allocated Emerging Technology Fund for foreign operations and related programs of USAID and the Department of State.

SEC. ____.—EMERGING TECHNOLOGY FUND.—

(a) ESTABLISHMENT.—There is established within the Department of State an Emerging Technology Fund (“Fund”) to facilitate holistic planning of digital foreign assistance, digital development projects, emerging technology programs, and other related initiatives of the Department of State and the United States Agency for International Development and to ensure the efficient management, coordination, operation, and utilization of such resources.

(b) FUNDING.—Funds otherwise available for the purposes of subsection (a) may be deposited in such Fund.

(c) AVAILABILITY.—Amounts deposited into the Fund shall remain available until expended.

(d) EXPENDITURES FROM FUND.—Amounts deposited in the Fund shall be available for the purposes of subsection (a).

(e) TRANSFER AUTHORITY.—Amounts available in the Fund may be transferred to any account of the Department of State or the United States Agency for International Development authorized by the Secretary of State for the purposes of carrying out a program described in subsection (a). Any amount so transferred shall be credited to the account to which it is transferred. The transfer authority provided in this subsection is in addition to any other transfer authority available to the Department of State.

Recommendation: Enhance the United States' Position as an International Digital Research Hub.

Component #2: Establish the Multilateral AI Research Institute (MAIRI).

SEC. ____.—MULTILATERAL ARTIFICIAL INTELLIGENCE RESEARCH INSTITUTE.—

(a) ESTABLISHMENT.—Not later than 180 days after the date of the enactment of this Act, the Director of the National Science Foundation (“Director”) shall establish a Multilateral Artificial Intelligence Research Institute (“MAIRI”) that leverages the National Artificial Intelligence Research Institutes as well as contributions from international partners, U.S. Government agencies, and non-governmental partners to facilitate international collaborative research and development initiatives involving artificial intelligence (“AI”). MAIRI shall have both a physical center located in the United States and a virtual presence.

(b) PURPOSE.—The purpose of MAIRI shall be to facilitate collaboration of international artificial intelligence research, foster international artificial intelligence innovation, and develop the next generation global artificial intelligence workforce in a manner that comports with democratic values and helps to preserve free and open societies.

(c) INTERNATIONAL PARTNERS.—As authorized by section 1872 of title 42, United States Code, the Director, in coordination with the Secretary of State, shall seek to develop partnerships with foreign governments that have existing research agreements and collaborative relationships with the United States. The Director of MAIRI shall provide for international partners to collaborate in the governance of MAIRI, contingent upon appropriate contributions of financial support.

(d) OTHER PARTNERS.—To further the goals of MAIRI, the Director shall seek, as necessary, partnerships with other U.S. Federal departments and agencies, and their national laboratories, and non-governmental partners, such as from industry, academia, research institutions, and philanthropies on a project-by-project basis.

(e) FACILITATION.—The Director, in coordination with the Secretary of State, shall facilitate the operations of MAIRI by creating a trusted learning cloud and associated compute capacity to facilitate international collaborative research by enabling access to needed resources, compute, and data for shared innovation, research, and development

(f) RESEARCH AGENDA.—MAIRI shall work with international partners, as well as U.S. Government partners, as needed, to—

(1) develop principles for multilateral artificial intelligence research, which address the importance of research integrity, the need for transparency, the necessity of open data and data sharing, the development of risk-benefit frameworks, and the use of merit-based competition reviews for research proposals; and

(2) develop research priorities that leverage members' capabilities and may include the development of—

(A) shared, secure compute resources, including joint benchmarking projects and data sharing, pooling, and storing initiatives founded on commonly agreed principles that ensure trust, privacy and security;

(B) privacy-preserving artificial intelligence and machine learning technologies, including technologies like federated learning and on-device prediction that enable remote execution, encrypted computation through multi-party computation and homomorphic encryption, and differential privacy; and

(C) smart city technologies, aligned with democratic values, that promote sustainability as well as norms that should guide standards development at bodies like the ITU and technical standards bodies.

(g) SOLICITATION AUTHORIZED.—The Director is authorized to issue one or more solicitations to create a physical facility to support the establishment of MAIRI. Any such solicitation shall provide for the selection of an awardee on a competitive, merit-reviewed basis.

(h) FINANCIAL ASSISTANCE TO ESTABLISH AND SUPPORT MAIRI.—Subject to the availability of funds appropriated for this purpose, the Director, the Secretary of Energy, the Secretary of State, the Secretary of Commerce, and other Federal agency heads may award financial assistance, as determined by an agency head, to establish and support MAIRI and associated research.

(i) AUTHORIZATION OF APPROPRIATION.—There is authorized to be appropriated for fiscal years 2022 through 2027, in such funds as may be required, for the purpose of—

(1) establishing and maintaining a physical center for MAIRI in the United States;

(2) carrying out MAIRI research initiatives in cooperation with the National Science Foundation, the Department of Energy, the Department of State, and other appropriate federal agencies;

(3) creating a trusted learning cloud and associated compute capacity to facilitate international collaborative research;

(4) U.S. researchers' travel and associated expenses to participate in MAIRI workshops, conferences, and similar events; and

(5) the establishment of an endowment fund in cooperation with international partners.

Recommendation: Reorient U.S. Foreign Policy and the Department of State for Great Power Competition in the Digital Age.

Expedite necessary reorganization of the Department of State by passing legislation to create an Under Secretary for Science, Research and Technology (Q).

SEC. ____.—UNDER SECRETARY OF STATE FOR SCIENCE, RESEARCH AND TECHNOLOGY.—

(a) POSITION ESTABLISHED.—Subsection (b) of section 2651a of title 22, United States Code, is amended—

(1) in paragraph (1), by striking “6” and inserting “7”;

(2) by redesignating paragraph (4) as paragraph (5); and

(3) by inserting before redesignated paragraph (5) the following new paragraph:

“(4) UNDER SECRETARY FOR SCIENCE, RESEARCH AND TECHNOLOGY. There shall be in the Department of State, among the Under Secretaries authorized by paragraph (1), an Under Secretary for Science, Research and Technology, who shall have primary responsibility to assist the Secretary and the Deputy Secretary on matters related to international science and technology policy.”

(b) REORGANIZATION REQUIRED.—Not later than 180 days after the date of the enactment of this Act, the Secretary of State shall develop a plan to consolidate the

science and technology policy functions of the Department in a single division under the leadership of the Under Secretary for Science, Research and Technology.

CHAPTER 16: ASSOCIATED TECHNOLOGIES

Blueprint for Action

Recommendation: Foster a Vibrant Domestic Quantum Fabrication Ecosystem.

Enact a package of provisions that incentivizes the domestic design and manufacturing of quantum computers and their constituent materials.

SEC. ____.—TAX CREDIT FOR DOMESTIC DESIGN AND MANUFACTURING OF QUANTUM COMPUTERS AND CONSTITUENT MATERIALS.—

Section 41(d) of title 26, United States Code, is amended by adding at the end a new paragraph (5), as follows—

“(5) SPECIAL RULE FOR DOMESTIC DESIGN AND MANUFACTURING OF QUANTUM COMPUTERS AND CONSTITUENT MATERIALS.—

“(A) With regard to domestic design and manufacturing of qualified quantum computers and constituent materials, the term ‘qualified research’ shall include, in addition to research described in paragraph (1)—

“(i) the development and production of qualified quantum computers and constituent materials in the United States; and

“(ii) the training of United States persons with regard to the development and production of qualified quantum computers and constituent materials.

“(B) In this paragraph, the term ‘qualified quantum computers and constituent materials’ means—

“(i) any computers have been identified by the Secretary, in consultation with the Secretary of Commerce, as quantum computers; and

“(ii) any components or constituent parts of such computers that have been identified by the Secretary, in consultation with the Secretary of Commerce, as critical to the operation of such computers.”

General Note: Should Congress establish a National Technology Foundation pursuant to the Commission’s Chapter 11 recommendation, Congress should also review conflicting National Science Foundation authorities and delegating appropriate authorities to the NTF.

Appendix E: Funding Recommendation Table

Funding Recommendation Table

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 1 – Emerging Threats in the AI Era	1	Create a Foreign Malign Influence Response Joint Interagency Task Force (JIATF).	Office of the Director of National Intelligence	\$30 million	-
	2	Increase DARPA funding for media authentication, disinformation detection, attribution, and disruption.	Department of Defense: USD(R&E) - DARPA	\$60 million to \$80 million	-
	3	Fund a machine speed AI-enabled cyber defense acceleration study.	Department of Homeland Security	\$10 million	-
	4	Increase DARPA funding for AI-enabled cyber defense research.	Department of Defense: USD(R&E) - DARPA	\$20 million	-
	5	Increase National Institute of Standards and Technology AI testbed funding.	National Institute of Standards and Technology	\$25 million	-
	6	Provide funding for a SolarWinds threat review.	Cyberspace Solarium Commission	\$6.5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 2 – Foundations of Future Defense	1	Establish a dedicated AI Fund.	Department of Defense: USD(R&E)	\$200 million	-
	2	Increase investments in AI R&D.	Department of Defense	\$8 billion	-
	3	Establish a fund to to accelerate procurement and integration of commercial AI solutions for business applications.	Department of Defense: Joint Artificial Intelligence Center	\$100 million	-
	4	Provide funding to build enterprise data sets.	Department of Defense: Office of the Chief Data Officer	\$125 million	-
	5	Provide funding for technology scouting tools, data, and a technology fellows program.	Department of Defense, USD(R&E)	\$10 million	-
Chapter 3 – AI and Warfare	1	Develop innovative operational concepts that integrate new warfighting capabilities with emerging technologies.	Department of Defense: USD(R&E)	\$5 million	-
	2	Incentivize experimentation with AI-enabled applications through the Warfighting Lab Incentive Fund (WLIF).	Department of Defense: USD(R&E)	\$10 million	-
	3	Encourage a culture of "Thinking Red."	Department of Defense: Joint Warfighting Analysis Center	\$2.5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 3 – AI and Warfare	4	Direct the military services, in coordination with the Under Secretary of Defense (for Acquisition and Sustainment), the Joint Staff, and the Defense Logistics Agency, and enabled by enterprise services and expertise at the JAIC, to prioritize integration of AI into logistics and sustainment systems wherever possible.	Department of Defense: Office of the Deputy Secretary of Defense	\$100 million	-
	5	Define a joint warfighting network architecture by the end of 2021.	“Department of Defense: Office of the Chief Information Officer”	\$5 million	-
Chapter 5 – AI and the Future of National Intelligence	1	Work with the intelligence community to establish a 10-year, \$1 billion, Program of Record to provide long-term, predictable funding for technologies identified in the technology annex to the National Intelligence Strategy.	Office of the Director of National Intelligence	\$1 billion annually for FYs 2022-2032	-
Chapter 6 – Technical Talent in Government	1	Congress should create a National Reserve Digital Corps.	Office of Management and Budget	\$16 million	-
	2	Congress should establish a STEM Corps.	Department of Defense	\$5 million for FY 2022 & \$5 million for FY 2023	-
	3	Congress should create a United States Digital Service Academy.	New Entity	\$40 million initial appropriation	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 7 – Establishing Justified Confidence in AI Systems	1 Appoint responsible AI leads and supporting staff in each agency critical to national security.	Department of Defense; Office of the Director of National Intelligence; Department of Homeland Security; Federal Bureau of Investigation; Department of State; Department of Energy; & Department of Health and Human Services	\$21.5 million	This funding supports one responsible AI lead and two supporting staff. Additionally, the funding includes responsible AI leads for each of the armed services in the Department of Defense and each of the agencies of the Intelligence Community.
Chapter 8 – Upholding Democratic Values	1 Congress should establish third-party testing center(s) to allow independent, third-party testing of national security-related AI systems that could impact U.S. persons.	National Institute of Standards and Technology	\$1.2 million	-
Chapter 9 – A Strategy for Competition and Cooperation	1 Create a Technology Competitiveness Council.	The White House: Executive Office of the President	\$2 million	-
Chapter 10 – The Talent Competition	1 Congress should pass a new National Defense Education Act.	Department of Education; National Science Foundation	One time appropriation of \$8.2 billion	-
Chapter 11 – Accelerating AI Innovation	1 Establish a National Technology Foundation.	New Entity	\$30 million initial appropriation for start-up expenses; \$1 billion for FY 2022; \$5 billion for FY 2023; \$10 billion for FY 2024; \$15 billion for FY 2025; & \$20 billion for FY 2026	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 11 – Accelerating AI Innovation	1	Increase federal funding of Non-Defense AI R&D at compounding levels.	Multiple agencies, including: the NSCAI proposed National Technology Foundation; National Science Foundation; Department of Energy; National Institute of Standards and Technology; National Institutes of Health; & National Aeronautical and Space Administration	\$2 billion for FY 2022; \$4 billion for FY 2023; \$8 billion for FY 2024; \$16 billion for FY 2025; & \$32 billion for FY 2026	-
	2	Expand the Network of AI Research Institutes.	National Science Foundation	\$200 million for FY 2022; \$200 million for FY 2023; & \$200 million for FY 2024	-
	3	Establish an AI Innovator Award.	National Science Foundation	\$125 million	-
	4	Establish a team-based AI Award.	National Science Foundation	\$50 million for FY 2022; \$100 million for FY 2023; \$150 million for FY 2024; \$200 million for FY 2025; & \$250 million annually for FYs 2026-2028	-
	5	Implement the NAIRR Roadmap.	National Science Foundation	\$30 million	-
	6	Fund an AI Data Program.	Department of Energy	\$25 million	-
	7	Sponsor an Open Knowledge Network.	National Science Foundation	\$25 million	-
	8	Form a network of Regional Innovation Clusters.	National Institute of Standards and Technology	\$200 million for FYs 2022-2026	Funding recommended at \$20 million per Regional Innovation Cluster

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 13 – Microelectronics	1	Increase federal grants for microelectronics manufacturing.	Department of Commerce	\$15 billion total	\$3 billion per project on average
	2	Increase funding for DARPA’s Electronics Resurgence Initiative (ERI).	Department of Defense: USD(R&E) - DARPA	\$400 million for FY 2022 & \$5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	3	Increase funding for National Science Foundation semiconductor research.	National Science Foundation	\$300 million for FY 2022 & \$2.5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	4	Increase funding for Department of Energy semiconductor research.	Department of Energy	\$400 million for FY 2022 & \$4.5 billion total for FYs 2022-2026	These funding levels should ramp up on an annual basis as absorptive capacity increases
	5	Establish the Advanced Packaging National Manufacturing Program.	National Institute of Standards and Technology	\$1 billion for FY 2022 & \$5 billion total for FYs 2022-2026	-
	6	Establish the National Semiconductor Technology Center.	Department of Commerce in collaboration with the Department of Defense and Department of Energy	\$100 million FY 2022 & \$2 billion total for FYs 2022-2026	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	1 Provide funding for U.S. International Development Finance Corporation to execute development financing for technology infrastructure projects.	U.S. International Development Finance Corporation	\$1 billion	-
	2 Provide funding to support U.S. International Development Finance Corporation development financing initiatives.	Department of State; U.S. Agency for International Development	\$200 million	-
	3 Provide funding for U.S. Agency for International Development Digital Strategy.	U.S. Agency for International Development: Bureau of Democracy, Development, and Innovation	\$200 million	-
	4 Provide funding for an Interagency AI Standards team to support National Institute of Standards and Technology AI Standards Coordinator and fund travel and other administrative needs.	National Institute of Standards and Technology; Department of Defense; Department of State; Office of the Director of National Intelligence; Department of Energy; Department of Homeland Security; U.S. Agency for International Development	\$3.3 million	Funding includes five full-time employee (FTE) from National Institute of Standards and Technology and one FTE from each of the following departments and agencies: Department of Defense, Department of State, Office of the Director of National Intelligence, Department of Energy, Department of Homeland Security, and U.S. Agency for International Development.
	5 Provide funding to support grants for small- and medium-sized businesses to participate in international data and technical standards efforts.	Small Business Administration	\$1 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	6 Funding for administrative costs associated with establishing an U.S. Center of Expertise relationship with GPAI/OECD.	National Science Foundation	\$1 million	-
	7 Funding for the Multilateral AI Research Initiative (MAIRI), including establishing and maintaining physical center; supporting research initiatives; created a trusted learning cloud resource; and supporting U.S. researchers' travel and involvement in workshops, conferences, and events.	National Science Foundation; Department of State; Department of Energy	\$12.15 million annually for FYs 2022-2027	\$10 million to National Science Foundation/ Department of State/ Department of Energy for research and personnel; \$2M to National Science Foundation for infrastructure; \$150,000 to National Science Foundation for administrative costs.
	8 Provide funding for trusted learning cloud to facilitate collaborative R&D with allies and partners (envisioned as a component of MAIRI).	National Science Foundation; Department of State	\$11.3 million	Funding includes underlying infrastructure, data storage and sharing capacity, grants for researchers, foreign assistance grants.
	9 Provide funding to support grants for scholars and researchers to participate in international data and technical standards efforts.	Department of State	\$5 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
Chapter 15 – A Favorable International AI Order	10 Provide funding for immediate augmentation and training of U.S. diplomatic corps for efforts related to AI and emerging technology (funding does not include future funding needs which we recommend be determined by a focused planning effort to be undertaken by Department of State).	Department of State	\$8 million	\$550,000 - STAS; \$550,000 - Office of Communication and Information Policy; \$400,000 - Office of Science and Technology Cooperation; \$3.8 million - Regional Technology Officers (12 locations); \$1.25 million - Office of the Special Representative to Silicon Valley; \$450,000 - FSI training.
	11 Provide funding for the Bureau of Cyberspace Security and Emerging Technologies.	Department of State	\$20 million	-
	12 Provide funding for public diplomacy and engagement activities on AI innovation and democratic values.	Department of State	\$5.5 million	-
	13 Provide funding for AI exchange programs to support U.S. values and fund participation by developing countries in multilateral AI activities.	Department of State	\$8.5 million	-
	14 Provide funding for efforts to promote U.S. innovation and values and support American Spaces, Tech Camps, Maker Spaces, Speakers Program, and other initiatives.	Department of State	\$3 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail
	15 Provide funding for tracking and analysis of public opinion to measure impact of engagement efforts and guide strategic planning.	Department of State	\$1 million	-
	16 Provide funding for U.S. Science Envoys and Embassy Science Fellows programs.	Department of State	\$1 million	-
Chapter 15 – A Favorable International AI Order	17 Provide funding to support U.S. leadership in AI through Emerging Technology Coalition and internal programs.	Department of State: Office of the Under Secretary for Economic Growth, Energy, and the Environment (E)	\$5.5 million	Funding includes ETC support, creation of an advisory committee on emerging technology, private sector engagement, multilateral R&D efforts, tech-oriented diplomatic efforts, innovation enhancements.
	18 “Funding to support promotion of human rights and fundamental freedoms in AI context through civil society initiatives, promoting AI and emerging tech to counter censorship, and supporting research and awareness campaigns”	Department of State: Office of the Under Secretary for Civilian Security, Democracy, and Human Rights (J): Bureau of Democracy, Human Rights, and Labor (DRL)	\$1.5 million	-
	19 Provide funding to support use of AI for national security/ military applications through cooperation with allies and partners, to include joint exercises, grants, fellowships, and other activities.	“Department of State: Office of the Under Secretary of State for Arms Control and International Security (T)”	\$3 million	-

Chapter	Recommendation	Cabinet Departments, Major Agencies, and Program Offices	Amount	Appropriations Detail	
Chapter 15 – A Favorable International AI Order	20	Provide funds to support building technical capacity in emerging democracies and market economies to counter malign influence.	Department of State	\$3 million	-
	21	Provide funds to support research grants on malign influence in AI ecosystems.	Department of State	\$2 million	-
	22	Provide funds to support public diplomacy initiatives on international AI standards and tracking and reporting of impact on public engagement.	Department of State	\$2 million	-
	23	Provide funds to support US Global Innovation through Science and Technology (GIST) Initiative.	Department of State	\$1 million	-
	24	Provide additional funding to support foreign assistance activities around emerging tech and digital infrastructure, to include planning, assessments, and provision of assistance. Funds would support targeted, digital programs in several areas, including rule of law (INL), democracy and human rights (DRL), security cooperation (AVC/PM/ISN), and technical assistance (EB, STAS, others).	Department of State	\$230 million	-

*Unless otherwise noted funding is annual beginning in Fiscal Year 2022.

**All funding figures should be considered initial estimates for consideration by Congress and the Executive Branch.

Appendix F:

Commissioner Bios



Dr. Eric Schmidt, Chair

Dr. Eric Schmidt is an accomplished technologist, entrepreneur, and philanthropist. He joined Google in 2001 and helped grow the company from a Silicon Valley startup to a global leader in technology alongside founders Sergey Brin and Larry Page. Schmidt served as Google's Chief Executive Officer and Chairman from 2001 to 2011, as well as Executive Chairman and Technical Advisor. Under his leadership, Google dramatically scaled its infrastructure and diversified its product offerings while maintaining a strong culture of innovation.

In 2017, he co-founded Schmidt Futures, a philanthropic initiative that bets early on exceptional people making the world better. Schmidt is the host of "Reimagine with Eric Schmidt," a podcast series of conversations with leaders to explore how society can build a brighter future after the COVID-19 pandemic.



The Honorable Robert Work, Vice Chair

Robert Work was the 32nd Deputy Secretary of Defense, serving alongside three Secretaries of Defense from May 2014 to July 2017. In 2001, he retired as a Colonel in the United States Marine Corps after spending 27 years on active duty. He subsequently served as Senior Fellow and Vice President and Director of Studies at the Center for Strategic and Budgetary Assessments. In January 2009, he was asked to join the Obama administration as the 31st Under Secretary of the Navy, and was confirmed in that role in May 2009. Work stepped

down as the Under Secretary in March 2013 to become the Chief Executive Officer for the Center for a New American Security (CNAS). He remained in that position until he assumed the role of Deputy Secretary of Defense in May 2014. He currently is the President and Owner of TeamWork, LLC, which specializes in defense strategy and policy, programming and budgeting, military-technical competitions, revolutions in war, and the future of war.



Safra Catz


Safra A. Catz has served as chief executive officer of Oracle Corporation since 2014 and a member of the company's board of directors since 2001. She joined Oracle in 1999 and held various positions within the company, including President and Chief Financial Officer, prior to being named CEO. Catz currently serves as a director of The Walt Disney Company and previously served as a director of HSBC Holdings plc.



Dr. Steve Chien

Dr. Steve Chien is a Technical Fellow, Senior Research Scientist, and the Technical Group Supervisor of the Artificial Intelligence Group at the Jet Propulsion Laboratory, California Institute of Technology. Chien has led the deployment of AI software to a wide range of missions. He is currently supporting the development of onboard and ground automated scheduling for the Mars 2020 rover mission, as well as scheduling technologies for the ECOSystem Spaceborne Thermal Radiometer Experiment on Space Station (ECOSTRESS) and Orbiting Carbon


Observatory 3 (OCO-3). Chien has received numerous awards for these efforts, to include Lew Allen Award for Excellence, JPL's highest award recognizing outstanding technical achievements by JPL personnel in the early years of their careers. He has been recognized four times in the NASA Software of the Year competition and has received four NASA medals for his work in AI for space. In 2011, Chien was awarded the inaugural American Institute of Aeronautics and Astronautics Intelligent Systems Award for his contributions to spacecraft autonomy.



The Honorable Mignon Clyburn

Mignon L. Clyburn served as Commissioner on the Federal Communications Commission (FCC) from 2009 to 2018, and acting chair from May to November of 2013. During her nearly nine years at the FCC, Mignon was committed to closing persistent digital and opportunities divides that continue to challenge rural, Native, and low wealth communities. Previously, Clyburn served for 11 years on the South Carolina Public Service Commission. Prior to that, she was the publisher and general manager of the Coastal Times,

a family-founded, Charleston-based weekly newspaper focusing on issues affecting the African American community. Clyburn is currently the principal of MLC Strategies, LLC.



Christopher Darby

Christopher Darby has served as President and CEO of In-Q-Tel since September 2006. He is also a member of its Board of Trustees. Prior to joining In-Q-Tel, Darby was a Vice President and General Manager at Intel, where he oversaw the Middleware Products Division. He joined Intel in August 2005 with the acquisition of Sarvega, a venture-backed supplier of XML networking and security products, where he served as President and CEO. Prior to Sarvega, Darby was the Chairman and CEO of @stake, an Internet security consulting firm ultimately acquired by

Symantec. Before that, Darby served as President and CEO of Interpath Communications, which was later acquired by US Internetworking. Earlier in his career, he held several executive positions at Digital Equipment Corporation (now Hewlett-Packard) and Northern Telecom (now Nortel Networks). Chris began his career at Bell Northern Research.



Dr. Kenneth Ford

Dr. Kenneth Ford is Founder and CEO of the Institute for Human & Machine Cognition. His research interests include AI, human-centered computing, and human performance and resilience. Ford is a Fellow of the Association for the Advancement of AI (AAAI), and a charter Fellow of the National Academy of Inventors. He has received many awards and honors including the Doctor Honoris Causas from the University of Bordeaux in 2005, the 2008 Robert Englemore Award for his work in AI, and the AAAI Distinguished Service Award in 2015. In 2015, he was elected as Fellow of

the American Association for the Advancement of Science and in 2017 was inducted into the Florida Inventors Hall of Fame. Ford has served on the National Science Board, the Air Force Science Advisory Board, and the Defense Science Board. In 2008, he was named as Chairman of the NASA Advisory Council – a capacity in which he served through 2011. In 2010, Ken was awarded NASA's Distinguished Public Service Medal – the highest honor the agency confers.



Dr. José-Marie Griffiths

Dr. José-Marie Griffiths is president of Dakota State University in Madison, South Dakota. Griffiths has spent her career in research, teaching, public service, corporate leadership, economic development, and higher education administration. She has served in presidential appointments to the National Science Board, the U.S. President's Information Technology Advisory Committee, and the U.S. National Commission on Libraries and Information. Griffiths has led projects for more than 28 U.S. federal agencies such as the National Science Foundation

and NASA, and more than 20 major corporations including, AT&T Bell Laboratories and IBM, in more than 35 countries. She also has worked with seven major international organizations, including NATO and the United Nations. She has received over 20 significant awards in science, technology, teaching, and the advancement of women in these fields.



Dr. Eric Horvitz

Dr. Eric Horvitz is a technical fellow at Microsoft, where he serves as the company's first Chief Scientific Officer. Horvitz provides cross-company leadership and perspectives on advances and trends on scientific matters, and on issues and opportunities arising at the intersection of technology, people, and society. He is recognized for his research on challenges and opportunities with the uses of AI technologies amidst the complexities of the open world. Horvitz is the recipient of the Feigenbaum Prize and the Allen Newell Prize for contributions to AI.



Andrew Jassy

Andy Jassy is the founder and CEO of Amazon Web Services (AWS), the world's most comprehensive and broadly adopted cloud platform. Jassy launched AWS in 2006 and has managed an inventive and nimble team that has delivered more than 165 services for compute, storage, networking, databases, analytics, mobile, Internet of Things, Artificial Intelligence, security, hybrid, and enterprise applications. Prior to founding AWS, Jassy held

several leadership positions across Amazon. Shortly after joining the company in 1997, he authored the business plan for Amazon's Music business and served as its Director of Product Management and General Manager. Jassy also started the Amazon Customer Relationship Management team, led marketing for Amazon, and was Technical Advisor (shadow) to Amazon Founder and CEO Jeff Bezos.



Gilman Louie

Gilman Louie is Co-Founder and Partner of Alsop Louie Partners, an early-stage technology venture capital firm founded in 2006. From 1999 until 2006, Louie was the first CEO of In-Q-Tel. Prior to In-Q-Tel, Louie built a career as a pioneer in the interactive entertainment industry, during which he founded and ran a publicly traded company called Spectrum HoloByte, and served as Chief Creative Officer of Hasbro Interactive. He serves as a member of the Board of Directors for the Markle Foundation, Maxar Technologies, Niantic, Lookingglass Cyber Solutions,

Aurora Insights and various other private companies and non-profit foundations. He is also Chairman of the Board of the Federation of American Scientists. Louie has served as a member of the Technical Advisory Group of the United States Senate Select Committee on Intelligence, and as a Commissioner of the National Commission for Review of Research and Development Programs of the United States Intelligence Community. He has received dozens of awards for his achievements, including from the NGA, CIA, and DNI, and in 2002 was named as one of fifty scientific visionaries by Scientific American.



Dr. William Mark

Dr. William Mark leads SRI International's Information and Computing Sciences division, creating new technology in machine learning, virtual personal assistance, trusted systems, and speech and vision analytics. The group also commercializes technology, licensing to corporations and creating spinoff companies such as Siri, Kasisto, CurieAI, and LatentAI. Prior to joining SRI International, Mark headed research groups at National Semiconductor, Lockheed Martin, and the University of Southern California Information Sciences Institute.



Dr. Jason Matheny

Dr. Jason Matheny is the founding director of Georgetown University's Center for Security and Emerging Technology (CSET). Previously he served as Assistant Director of National Intelligence, and Director of IARPA, responsible for the development of breakthrough technologies for the U.S. intelligence community. Before IARPA, he worked at Oxford University, the World Bank, the Johns Hopkins University Applied Physics Laboratory, the Center for Biosecurity, and Princeton University, and was the co-founder of two biotechnology companies.



The Honorable Katharina McFarland

Katharina McFarland serves as Chairman of the Board of Army Research and Development at the National Academies of Science, and as a Director on the Boards of SAIC, Exyn Technologies, and the Procurement Round Table. With more than 30 years of government service, McFarland is widely recognized as a leading subject matter expert on government procurement. She also serves as an advisor to Raytheon Missile Systems Division Senior Advisory Board, Cypress International Senior Strategy Group, Transunion Corporation Advisory

Board, and Sehlke, Inc. Senior Advisory Board. From 2012 to 2017, McFarland served as the Assistant Secretary of Defense for Acquisition and as acting Assistant Secretary of the Army (Acquisition, Logistics & Technology) from 2016-2017. She was President of the Defense Acquisition University from 2010 to 2012, and the Director of Acquisition at the Missile Defense Agency from 2006 to 2010. She has received an Honorary Doctoral of Engineering from the University of Cranfield in the United Kingdom, the Presidential Meritorious Executive Rank Award, the Secretary of Defense Medal for Meritorious Civilian Service Award, the Department of the Navy Civilian Tester of the Year Award, and the Navy and United States Marine Corps Commendation Medal for Meritorious Civilian Service.



Dr. Andrew Moore

Dr. Andrew W. Moore is a distinguished computer scientist with expertise in machine learning and robotics. He became the head of Google Cloud Artificial Intelligence division in January 2019. Moore previously worked at Google from 2006 to 2014 and was the founding director of Google's Pittsburgh engineering office in 2006. He then spent a four-year hiatus at Carnegie Mellon University as the dean of the School of Computer Science. Moore's research interests encompass the field of "big data" — applying statistical methods and mathematical

formulas to massive quantities of information, ranging from web searches to astronomy to medical records, in order to identify patterns and extract meaning from that information. His past research has included improving the ability of robots and other automated systems to sense the world around them and respond appropriately.

Appendix G: Commission Staff and Contributors

EXECUTIVE STAFF

Yll Bajraktari,
Executive Director

Michael J. Lueptow,
General Counsel

Michael L. Gable,
Chief of Staff

Tara M. Rigler,
Director of Strategy,
Communications & Engagements

Angela A. Ponmakha,
Director of Operations,
Designated Federal Officer

LEGISLATIVE AFFAIRS

Brandon McKee

Jenilee Keefe Singer

SENIOR ADVISORS

Dr. Seth Center

Robert Nelson

RESEARCH AND ANALYSIS

Courtney Barno
Dr. Ryan Carpenter
Matthew Cordova
Caroline Danauy
Raina Davis
Tess deBlanc-Knowles
Rama Elluru
Michael Garris

Matthew Gentzel
Charles Howell
LTC Michael Jackson, USA
Rebekah Kennel
Jeffrey Kojac
David Kumashiro
CAPT Lance Lantier, USN
Christie Lawrence

Paul Lekas
Dr. Margaret Lentz
Quinn Lorenz
Justin Lynch
Col Paul “P.J.” Maykish, USAF
Kevin McGinnis
Christopher McGuire

Paul Rhodes
Dr. Christopher Rice
Joe Wang
Parker Wild
Jessica Young
Olivia Zetter

OPERATIONS AND LEGAL TEAM

Chelsea Holt
Sarah Johnson
Brent Myles

Jennifer Sheehan
Angela Stacks
Jamie Tomberlin

INTERNS

Richard Altieri
Madeline Blanchard
Sabrina Broderick
Shaantam Chawla
Devin Davidson
Nickie Deahl
Hudson Dizon
Dylan Halpern
Courtney Lange
Alexander Mann
Nikhil Marda
M. Marin Ruelas Mendoza

Ariana Orne
Sultan Seraj
Katie Stolarczyk
Jaide Tarwid
Christopher Tonelli
Claire Trotter
Samuel Trotter
Aristotle Vainikos
Jackson Valen
Zoe Weinberg
Kate Yeager


The Commission would like to thank Members of Congress, Congressional Staff, government personnel, industry professionals, academia, members of the public, and all others who participated in, advised on, or commented on our work. The unified effort of everyone involved made this document possible. Additionally, the Commission thanks all of our outside contributors whose hard work resulted in this product.

OUTSIDE CONTRIBUTORS


John Bansemer
Susanna Blume
Dr. Anne Bowser
Scott Britt
Kristy Colbert
Mark Cohen
Dr. David Danks
Jeffrey Ding
Dr. Kathleen Fisher
Amanda Foley
Dr. Kristin Gilkes
Dr. Bryce Goodman
Gregory Grant
Erin Hahn
Orin Hoffman
Dajonte Holsey
Dr. Michael Horowitz
Dr. Andrew Imbrie
Taylor Lineberger
Dr. Albana Shehaj
Dr. Paul Scharre
Dr. William Scherlis
Raj Shah
John “Jack” Shanahan
Dr. Bernadette Johnson

Elsa Kania
Dr. Christopher Kirchhoff
Zachary Kuehn
Thomas Kalil
Peter Levine
Frank Long
Michael “Brendan” McCord
Michael McNerney
Tariq Mehmood
Paul Michel
Dr. Nadia Schadlow Murphy
Adam Mossoff
Geoffrey Odlum
Scott Padgett
Jared C. Ponmakha
Douglas Rand
Dr. Heather Roff
Craig Smith
Michael Soos
Dean Souleles
Dr. Barbara Stephenson
Francoise von Trapp
German Wegbrait
Darren Wright
Dr. Amy Zegart

A special thanks to Lirijon Kadriu for designing the Commission’s logo.



NATIONAL
SECURITY
COMMISSION
ON ARTIFICIAL
INTELLIGENCE



**THE NATIONAL SECURITY COMMISSION
ON ARTIFICIAL INTELLIGENCE**