



Standing on the Shoulders of AI Giants

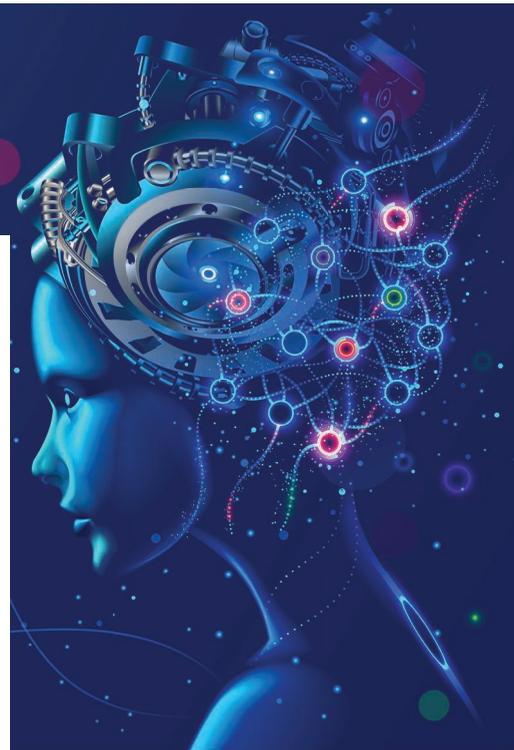
Hsiao-Ying Lin^{1b}, Huawei France

This article reviews the key artificial intelligence (AI) breakthroughs made by two AI innovation giants, OpenAI and DeepMind.

Technological innovations in artificial intelligence (AI) are like magic. Generative AI technology attempts to capture the creativity of human beings; for example, the Generative Pre-Trained Transformer 3 (GPT-3) model can generate written articles. AI-based decision-making algorithms accelerate the search process and can outperform humans; for example, AlphaGo outperforms professional Go players at the game Go. With a focus on two AI giants, OpenAI and DeepMind, we review some of the technological breakthroughs in AI of the last decade.

GIANT MINDS THINK ALIKE

OpenAI, founded in 2015, is an AI research and development company based in the United States. The company develops general AI technology to benefit all of



humanity. OpenAI began as a nonprofit company and later split into the “capped-profit” organization OpenAI LP and the parent nonprofit organization OpenAI Inc. in 2019. OpenAI has received high attention for its generative AI models, such as GPT-3 [1] and DALL-E (a portmanteau of

the artist Salvador Dalí and the Disney Pixar character, WALL-E) [2].

DeepMind is a U.K.-based AI company founded in 2010 that joined Google in 2014. Since 2015, DeepMind has been a wholly owned subsidiary of Alphabet Inc. DeepMind achieved early success in computer games, such as AlphaGo [3], and continues to solve fundamental scientific problems, such as with its development of AlphaTensor [4], which finds ways of solving the matrix multiplication problem with fewer multiplication operations. The company’s goal is to understand the mechanisms of the human brain and develop artificial general intelligence.

OpenAI and DeepMind have similar visions for implementing artificial general intelligence and share similar initial technological approaches for applying reinforcement learning. Because artificial general intelligence is the greatest challenge in AI, this shared vision is unsurprising. However, both companies also began by developing reinforcement learning approaches in gaming

environments. Reinforcement learning mimics how humans learn from trial-and-error experiences. The key challenges of applying reinforcement learning include properly modeling a

- › OpenAI Gym [5] is a platform for developing and benchmarking reinforcement learning algorithms in various gaming environments, including Atari

dynamics. In simple cases, the success rate was 60%. Although the AI has not been perfected, the Rubik's Cube program demonstrates how reinforcement learning can affect the physical world.

With a focus on two AI giants, OpenAI and DeepMind, we review some of the technological breakthroughs in AI of the last decade.

task as a learning problem and aligning the algorithms with the intended purpose. Soon after, OpenAI and DeepMind made great advancements in various AI domains. OpenAI has advanced generative technology, and DeepMind continues to solve fundamental scientific problems. Figure 1 summarizes some of the milestones and breakthroughs of these companies.

OPENAI ADVANCEMENTS

OpenAI began by providing friendly development environments called *OpenAI Gym* and *Universe* for designing and benchmarking reinforcement learning algorithms for various problems. The company also developed a Rubik's Cube program to demonstrate the real-world applications of reinforcement learning.

games and MuJoCo. By using OpenAI Gym, the performance of reinforcement learning algorithms in different environments can be evaluated and compared with that of other algorithms. Universe [6] is another software platform that converts software programs into environments for reinforcement learning tasks. With Universe, OpenAI Gym can support more environments, such as browser tasks and complex games.

- › By using their AI software and a robotic hand, OpenAI was able to successfully solve a Rubik's Cube. One challenge was properly modeling the physical world in the simulation to capture accurate physical

After OpenAI was restructured, most AI advancements were made in generative AI for text, images, and audio.

- › GPT-3, proposed in May 2020, is a generative model of natural language processing with 175 billion parameters. The model can write articles, summarize books, and autocomplete formatted text, and it enables various downstream tasks. One application of GPT-3 is OpenAI Codex, which translates natural language into programming code. OpenAI Codex was proposed in August 2021 and provides a natural language interface for developing software applications.
- › DALL-E, introduced in January 2021, is a GPT-3-based multi-modal model used for generating images from a given text prompt. The waitlist for DALL-E was removed in September 2022,

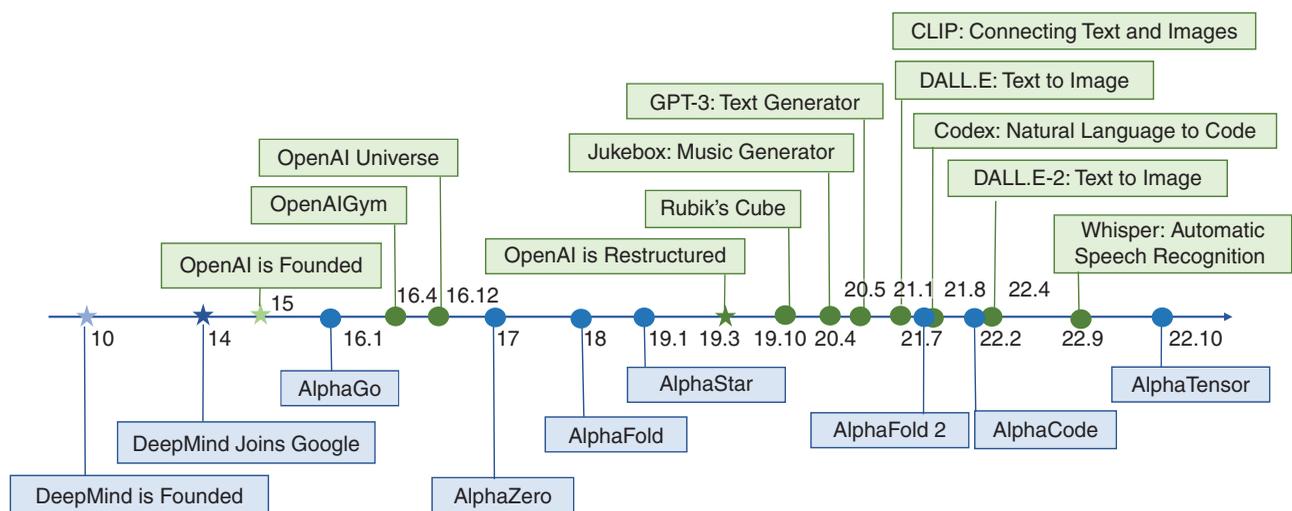


FIGURE 1. Milestones of OpenAI and DeepMind shown in green and blue, respectively, with time labels.

and more applications of DALL-E are expected. Figure 2 shows examples generated by DALL-E from the prompt “standing on the shoulders of giants.”

Although none of the images show a giant, they provide a similar conceptual idea. Figure 3 shows another set of examples from the prompt “giants in a rose garden,” where some giants appear. DALL-E 2 [7], proposed in April 2022, is the second generation of DALL-E and can combine concepts, attributes, and styles to generate images at higher resolutions. DALL-E 2 can also expand and edit images according to text descriptions. DALL-E 2 is a two-stage model. In the first stage, the model generates an embedded image from a text caption; in the second stage, a decoder generates an image based on the embedded image. Using diffusion models

to implement both stages yields high-quality images efficiently.

- › Jukebox [8] is a generative model for music. Given an input genre,

language. Moreover, Whisper can transcribe this recognized speech in English. Whisper was trained on 680,000 supervised

OpenAI has advanced generative technology, and DeepMind continues to solve fundamental scientific problems.

artist, and lyrics, the model outputs a new music sample. Jukebox tries to capture human creativity based on the chosen genre and artist.

- › Whisper [9] is a multimodal generative model proposed in September 2022. Given speech audio as input (not limited to English), Whisper is able to generate English transcripts. Whisper automatically recognizes input audio as speech and accounts for talking speed, accent, and

hours of multilingual web-collected audio data and uses encoder–decoder transformer model structures with five different model sizes for various scenarios.

DEEPMIND ADVANCEMENTS

DeepMind also began by exploring reinforcement learning in gaming environments to develop artificial general intelligence and created programs such as AlphaGo, AlphaZero, and AlphaStar.

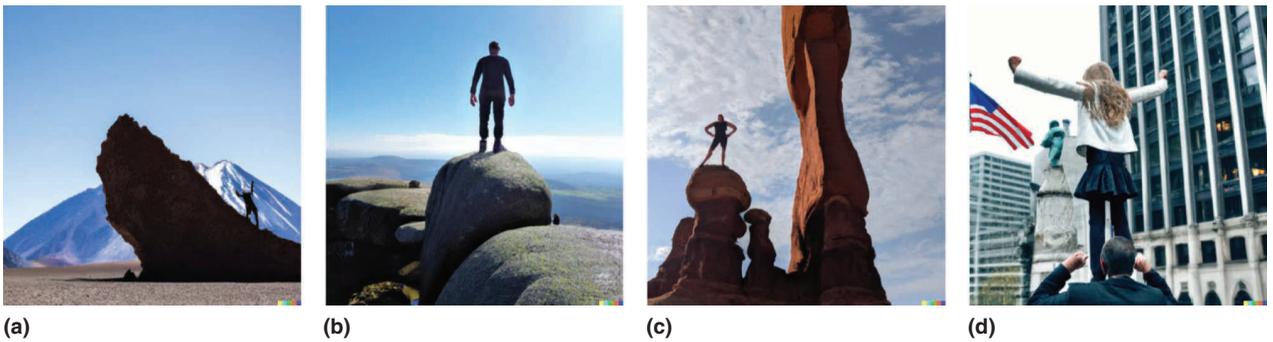


FIGURE 2. (a)–(d) Images generated by DALL-E from the input “standing on the shoulders of giants.”



FIGURE 3. (a)–(d) Images generated by DALL-E from the input “giants in a rose garden.”

- › AlphaGo [3], introduced in 2016, is famous for being the first computer program to defeat Lee Sedol, a professional Go player. Go is an ancient Chinese board game that requires strategic thinking. AlphaGo uses reinforcement learning and learns from playing with professional players. Lee Sedol once commented, “I thought AlphaGo was based on probability calculation and that it was merely a machine. But when I saw this move, I changed my mind. Surely, AlphaGo is creative. This move was really creative and beautiful.” Later, AlphaGo Zero was developed to test another learning approach, namely playing against itself and using initial strategies at random. AlphaGo Zero’s success indicates that input of human experience is not mandatory to train AI algorithms.
- › AlphaZero [10], introduced in late 2017, is a system that teaches itself how to master chess, shogi (Japanese chess), and Go. It extends the success of “self-play” strategy from Go to chess and shogi.
- › AlphaStar [11], introduced in 2019, is a reinforcement learning agent for *StarCraft II*, a multiplayer real-time strategy game. AlphaStar achieved the rank of Grandmaster, meaning that AlphaStar’s performance was in the top 0.4% (the best performance of AlphaStar made the top 0.15%). A new approach, multi-agent reinforcement learning, was developed to train AlphaStar. Since no one strategy can win in all circumstances, this learning

approach develops diverse strategies and later sophisticatedly integrates strategies.

In addition to exploring learning strategies from reinforcement learning in gaming environments, DeepMind applies derived principles and insights to solve fundamental scientific problems.

- › AlphaFold [12] provided a solution to the 50-year-old protein folding problem of biology, that is, predicting the 3D structure of proteins on the basis of a given sequence of amino acids. The shape of a protein is linked with its function and can hypothetically be determined from its sequence of amino acids. With greater knowledge of the 3D structures of proteins, biologists can conduct more advanced research. AlphaFold is quicker than conventional methods based on nuclear magnetic resonance, X-ray crystallography, or cryogenic electron microscopy. AlphaFold 2 [13], introduced in 2022, expanded the protein structure database from 1 million structures to more than 200 million structures.
- › AlphaCode [14] generates code to solve unseen programming problems. AlphaCode was created to advance the problem-solving skills of AI. On Codeforces, a website that hosts programming competitions, AlphaCode got an estimated rank within the top 54.3% of more than 5,000 participants. In Codeforces competitions, each problem has a description, input sample, and correct output sample. If the submitted program generates the correct output, the problem is solved. For each problem, 10 chances are allowed. AlphaCode uses a large-scale transformer-based language model to generate potential solutions and filters

them to the 10 most favorable candidates. The performance of AlphaCode demonstrates the problem-solving potential of AI provided with critical thinking, logic, algorithms, coding, and natural language knowledge.

- › AlphaTensor [4] provides a groundbreaking method of matrix multiplication. Matrix multiplication is a fundamental computing operation in various applications. Because multiplication operations are more computationally expensive than addition operations, researchers have sought a computational sequence for matrix multiplication with fewer multiplication operations (also called a decomposition technique). For example, computing a 4×5 matrix by a 5×5 matrix originally required 100 multiplication operations. This was initially improved to 80 operations and eventually to 76 operations by AlphaTensor. Thus, AlphaTensor can discover new computational sequences for matrix multiplication that require fewer multiplication operations. Moreover, for indicated specialized computing hardware, AlphaTensor can also find optimized computational sequences with lower cost.

This is a time full of miracles. Both OpenAI and DeepMind have made tremendous advancements in AI technology; their programs can create content, self-learn, and solve real-world problems. The future of AI is full of possibilities, and the progress made by current AI giants will continue to inspire and excite future generations. I cannot wait to see more AI innovations! 

REFERENCES

1. T. B. Brown et al., “Language models are few-shot learners,” 2021, *arXiv:2005.14165*.

DISCLAIMER

This article contains the views of the author. The opinions expressed here are hers alone.

2. A. Ramesh et al., "Zero-shot text-to-image generation," 2021, *arXiv:2102.12092*.
3. D. Silver et al., "Mastering the game of Go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, 2016, doi: 10.1038/nature16961.
4. A. Fawzi et al., "Discovering faster matrix multiplication algorithms with reinforcement learning," *Nature*, vol. 610, no. 7930, pp. 47–53, 2022, doi: 10.1038/s41586-022-05172-4.
5. G. Brockman et al., "OpenAI gym," 2016, *arXiv:1606.01540*.
6. "Universe." OpenAI. Accessed: Oct. 12, 2022. [Online]. Available: <https://openai.com/blog/universe/>
7. A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, "Hierarchical text-conditional image generation with CLIP latents," 2022, *arXiv:2204.06125*.
8. P. Dhariwal, H. Jun, C. Payne, J. W. Kim, A. Radford, and I. Sutskever, "Jukebox: A generative model for music," 2020, *arXiv:2005.00341*.
9. A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision." OpenAI. Accessed: Oct. 12, 2022. [Online]. Available: <https://cdn.openai.com/papers/whisper.pdf>
10. D. Silver et al., "A general reinforcement learning algorithm that masters chess, shogi and Go through self-play," *Science*, vol. 362, no. 6419, pp. 1140–1144, Dec. 7, 2018, doi: 10.1126/science.aar6404.
11. O. Vinyals et al., "Grandmaster level in StarCraft II using multi-agent reinforcement learning," *Nature*, vol. 575, no. 7782, pp. 350–354, 2019, doi: 10.1038/s41586-019-1724-z.
12. A. W. Senior et al., "Improved protein structure prediction using potentials from deep learning," *Nature*, vol. 577, no. 7792, pp. 706–710, 2020, doi: 10.1038/s41586-019-1923-7.
13. J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold," *Nature*, vol. 596, no. 7873, pp. 583–589, 2021, doi: 10.1038/s41586-021-03819-2.
14. Y. Li et al., "Competition-level code generation with AlphaCode," 2022, *arXiv:2203.07814*.

HSIAO-YING LIN is a principal researcher at Huawei France, 92100 Boulogne-Billancourt, France. Contact her at hsiaoying.lin@gmail.com.



IEEE Computer Graphics and Applications bridges the theory and practice of computer graphics. Subscribe to CG&A and

- stay current on the latest tools and applications and gain invaluable practical and research knowledge,
- discover cutting-edge applications and learn more about the latest techniques, and
- benefit from CG&A's active and connected editorial board.

