Your business guide:

AI Unwrapped

- Understanding AI technologies
- Mapping the next stage of your enterprise AI journey

EVIDEN

CONTENT

Foreword

AI has the potential to revolutionize what humans can do, from diagnosing, treating and preventing disease, to predicting and pre-empting sophisticated cyber-attacks, to making the vision of safe autonomous vehicles a reality.

Since the start of this decade, its use has exploded. With the more eye-catching capabilities of tools such as ChatGPT, the power of AI is also growing behind the scenes, shaping the way people work and live every day.

Today's enterprises know that if they don't invest now, they risk falling behind.

But they're also telling us that with this speed come some important questions. What's the difference between traditional AI and generative AI? How and where are these technologies best used? And where's the best place to start?

In parallel, new challenges are materializing around security, privacy, regulation and ethics. Where are the risks? Are AI-generated artefacts trustworthy? And what about legal compliance?

That's what this e-book explores. And it sets out six priorities for enterprises to consider in 2024 as your journey continues.

What's happened and why now?

2023 marked a threshold in the rapid evolution of AI. This was fuelled by years of growth in computing power, the availability of vast amounts of data from the internet , and the need to manage the complexity of an increasingly interconnected digital world (billions of users, soon trillions of Internet of Things (IoT) objects).

The most high-profile story was the race for generative AI between OpenAI (the creator of ChatGPT) and Microsoft, Google, Amazon, Meta, and pure GenAI players such as Anthropic, Mistral and others. In the media, this has overshadowed other novel AI technologies. One such advance is the emergence of 'liquid machine learning', which can learn from real-time and constantly changing data streams. This heralds a future in which intelligent systems will self-adapt to their environment, just as human beings do. And many other innovations are being conceived right now behind the doors of AI research laboratories.

Leaps of progress have also, of course, been made in other technological domains: the emergence of pioneering Quantum-as-a-service capabilities; the explosion of game-changing 5G IoT applications in multiple sectors; and spatial computing technologies – given the intensifying battle for future domination of the Metaverse, with Apple entering the arena.

Yet it's GenAI that has undoubtedly dominated the headlines.

First, some definitions

What is artificial intelligence?

Artificial intelligence (AI) is an umbrella term. It refers to use of computers and machines to simulate human intelligence and behaviours, for example solving complex problems and executing tasks to achieve the required results much faster and more accurately than humans.

What is machine learning?

Machine learning is a form of artificial intelligence that can analyze huge volumes of data to identify patterns in order to make decisions and predictions.

It uses self-learning algorithms to do this, which are trained using very large datasets, called training data.

In this way, a machine is trained in how to provide accurate results and perform specified tasks. The more data it learns from, the better it performs. Machine learning can scan vast data lakes, such as research data and social media platforms, to generate actionable insights.

Machine learning has a fast growing range of applications: rapid processing, facial recognition capabilities, sentiment analysis, cybersecurity (predicting, identifying and responding to cyber-threats in real time).


What is generative AI (GenAI)?

GenAI builds on machine learning as its foundation. Rather than just identifying patterns, it is focused on generating content, including text, sounds, images, video and even designs/structures or strategies.

While early GenAI technologies date back over decades, very rapid progress has been made since 2017. This is largely thanks to the invention of the 'transformer' technology, the effect of which is to dramatically accelerate and enrich the way machines can analyze relationships within sequential data such as text. Since then, an array of publicly available GenAI tools have taken hold – not just ChatGPT-4, but also Google Gemini, Amazon Bedrock, Meta Llama, and models from startups such as Hugging Face, Cohere, Anthropic or Mistral AI.

What defines GenAI is that it can both generate artefacts and also review and refine its outputs. It uses large language models and algorithms to analyze patterns in datasets to mirror the style or structure of specific types of content. There are four overarching categories of GenAI use cases.

- **Conversational user assistance to provide tailored responses based on context, information shared, and details requested:** chat interfaces & virtual assistants respond to user inquiries and customer service (such as information search or personalized user support). It dramatically expands on previous chatbots in a much more creative, immersive and hyper-personalized way.

- **Content ideation, creation, curation or aggregation (text, image, sound, video, synthetic data)** building creative content based on topics, inputs, style: automatically draft / synthesize / correct / translate a report, an image, a video, and so on.

- **Software development & delivery:** code creation and programming, test creation, business process, and task automation.

- **Strategy / product ideation or creation:** such as engineering parts, drugs or material design, architecture design, and strategies ideation.

While GenAI will dramatically change how people work, it should augment but not replace the role of people in originating work and ideas.
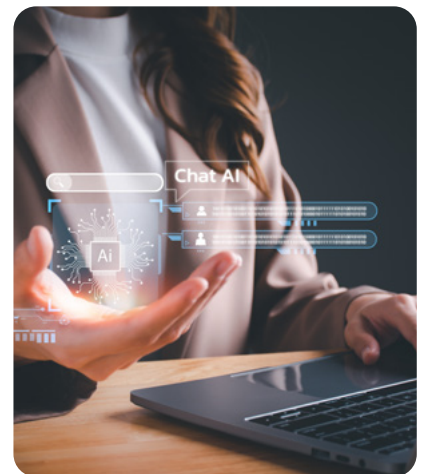
[1]In a nutshell, data volumes are doubling every three to four years –from the megabytes of the mainframe era, to the gigabytes of the network computing era, to terabytes in the age of social, mobile, analytics and cloud. Data volumes are already at 120 zettabytes. In 2030, we'll see the yottabyte. And the brontobyte is on the horizon for 2040.

Reimagining the world with AI

AI will fundamentally change the human-machine relationship. Machines will be much more than our tools: they will become our team-mates, mentors, friends, and even our customers or supervisors.

This will have dramatic consequences for customer interactions, operations, concepts of trust, and business models. The hype around GenAI is the tip of the iceberg. The future is about a much more holistic transformation that blends different evolving AI technologies to create enterprise applications.

So how is the rise of GenAI likely to drive this next phase of the AI revolution?

1. It will dramatically change productivity.

Traditional AI has already made process automation more intelligent. With GenAI comes the next step: intelligent knowledge management. This will transform the workplace, with most independent software vendors embedding GenAI into their applications suites. These are primarily administration, communication, collaboration, document creation/management and reporting suites (Copilot in MS Office, Gemini within Google Workplace, Einstein GPT in SalesForce, AI capabilities in enterprise resource planning (ERP) and customer relationship management (CRM) solutions from SAP and Oracle, and so on).

However, the more significant differentiators and reinventions will come when enterprises start to

integrate GenAI with their own data and workflows, either by using off-the-shelf GenAI models from GenAI vendors (trained on their data), or fine-turning models for their specific context, or even creating their own GenAI models. For example, Eviden has developed a 'Knowledge Pilot' GenAI tool that helps companies to transform internal knowledge management by reducing search costs, boosting productivity, and enhancing knowledge utilization in all fields related to their specific business. It is being deployed at several world-class corporations.

2. It will transform citizen or customer experiences.

For over a decade, big data technologies have dramatically accelerated and refined customer behavior analysis, segmentation, and next best-action analysis. With GenAI comes the next advance: automated creation of personalized content for each customer, from online conversations to personalized campaigns, ads and sales proposals. This is nascent. Most enterprises are still cautious in leveraging GenAI to shape and deliver customer experience because the technology is still emerging and there may be issues around the trustworthiness of the outputs. Yet

breakthroughs are happening. Eviden is working with Google, for example, to help a leading beauty company leverage GenAI to personalize the experiences of dozens of millions of customers across 40 countries. This is just one of many similar projects at Eviden in collaboration with hyperscalers and using various GenAI technologies.

3. It may ultimately revolutionize entire business models.

GenAI will be used to accelerate the creation of new disruptive capabilities and services (such as engineering or product design processes), to accelerate digital development (code creation), and to analyze and optimize business strategies. For the Julich Supercomputing Center, for instance, Eviden is building Europe's first Exascale supercomputer and probably the largest AI and GenAI supercomputer in the world. This is set to drive a paradigm shift in research and business innovation. By 2025, this unparalleled computing power will enable scientists and businesses to achieve what was previously impossible, in areas such as drug discovery, materials science and climate modeling.

The emergence of intelligent agents that incorporate both traditional AI and GenAI may be particularly disruptive. Imagine new enterprises leveraging swarms of intelligent agents, all collaborating to complete tasks in a highly effective and creative way. In other words, a vision of businesses that are autonomously developed and operated at near-zero marginal cost. While this may sound ambitious, the future often arrives sooner than expected. Early experimentations are already happening, paving the way for radical 'autonomous business' concepts.

Innovation will give a boost to all types of AI, with enterprises leveraging the best mix of solutions in each context – redefining how organizations operate in each industry, with multiple use cases including:
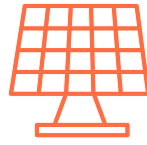
Energy and utilities: energy generation optimization, grid management, 360° prosumer services

Financial services: frictionless experience, smart operations, trading strategies, anti-fraud and anti-money-laundering
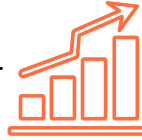
Healthcare and life sciences: patient care, drug discovery, precision medicine

Manufacturing: advanced engineering, smart factory, connected products

Public sector: citizen-centric services, public safety and defense, administrative automation, sustainability

Retail, transport and logistics: ubiquitous customer engagement, merchandising intelligence, logistics optimization

Telecoms and media: network optimization, targeted advertising, B2B services innovation, media content creation.

Ultimately, AI will accelerate the digital transformation of industry ecosystems, with dramatic impacts on value chains across ever sector.
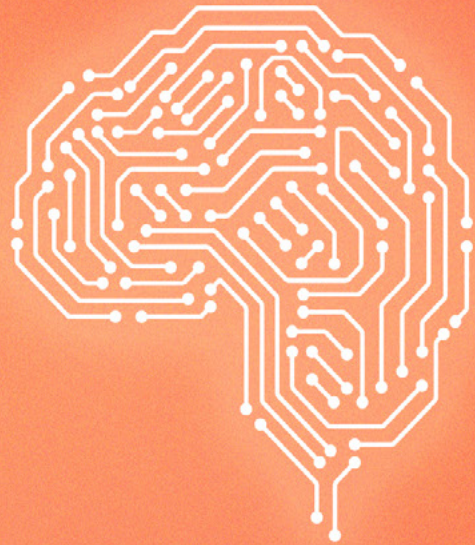
Where are the challenges and priorities for enterprises in 2024?

While progress has been startling, industry sectors are still at the start of this AI journey.

As we have seen during previous technological revolutions, the most agile organizations, not necessarily the biggest or those with most capital, may thrive in the age of AI.

What's very different, however, is the magnitude and speed of the advance. And the nature of AI and GenAI (the way it proliferates) means that any issues can be very rapidly replicated, amplified and multiplied.

Integration.

Integration of GenAI into legacy systems is a challenge, with more complexity than many enterprises expect. Any issues or siloes in the data will be reflected in the results that can be achieved. This complexity is compounded when GenAI vendors require data conversion into specific formats or dedicated infrastructures.

Trust, fairness and transparency.

Enterprises are often – and rightly – concerned about closed-source models, where the internal workings remain opaque. Vendor models sometimes behave unpredictably, either due to model upgrades or the nature of the training data, leading to inappropriate results (such as 'hallucinations'). Guardrails or filters provided by vendors are only partially effective, as they can't cover all risks and users often find creative prompts to circumvent them. Enterprises need to be prepared for people asking for unexpected things, maliciously (hackers often find new loopholes and win) or otherwise. This is one of the key reasons why enterprises are increasingly asking for open – or even open source – models, that can be further controlled and fine-tuned.

Cybersecurity.

GenAI doesn't necessarily change criminal techniques, but it enables large-scale operations and more convincing deceptions. For example, banks are alarmed at the growth of deepfake voice or video-based attacks.

Privacy, copyright and sovereignty. The potential impacts of evolving regulatory initiatives from China, the EU, the US, and others are still uncertain. New laws and policies are in development, including the launch of the UN AI Advisory Body, the signing of the AI Executive Order in the US, and agreement of the EU AI Act – all in late 2023. Sovereignty concerns, like China's refusal to use US technology and vice versa, are complicating international business operations (for example, which smart factory AI technologies are viable for plants in both the USA and China?).

Six priorities to explore

1. **Be bold.** Don't be afraid to be ambitious. What do you want to achieve? Are there exciting opportunities out on the horizon? The speed of change is accelerating, so having a north star to guide you not only keeps you on track, but could help you enrol stakeholders and achieve breakthroughs faster.

2. **Data, data, data.** Reassess and optimize your data strategies. How can you best expand, enrich, govern, protect and even monetize data? AI is only as good as the data that it's trained on. Making data AI-ready is critical because it is data that can give you sustainable advantage. Being AI-ready requires modern architectures, easing data access with Data Fabric (for unified access) and Data Mesh (for unified governance) strategies. Data Lakehouse technologies are often a good entry point to start the journey, within the context of a potentially large ecosystem.

3. **Responsible AI strategy.** 'Responsible AI' is multi-disciplinary and touches business, technology and legal domains. It's about understanding and mitigating any risks associated with using AI, rather than seeing it as a 'blackbox' in the IT department. Responsible AI strategies address the transparency and explainability of how the AI works and what its outputs are based on; fairness and bias mitigation to recognize and minimize any predispositions or partialities that arise; accountability for the effects that AI may have on people and society; and privacy and security to prevent leakage of confidential or sensitive information. Engaging and educating stakeholders (from the board, to application developers, to end users) is key to developing a culture of responsibility. Expert legal perspectives are needed to implement what's required for compliance of policies, processes, systems and data.

4. **Accelerate value-driven experimentation,** focusing on business rather than technology outcomes, with expertise so you can leverage for each use case the most appropriate mix of technologies to optimize the outcome/cost/performance ratio. The cost of AI, particularly GenAI, can be a concern. Large models training and inference requires huge computing power, and are therefore costly. Does your enterprise always need the most powerful model for each application? Horizon scanning all possibilities while focusing primarily on business outcomes will be essential.

5. **Integrate AI into your own processes and IP.** While all software vendors and technology service providers will embed AI technologies into their solutions and services, what will make the difference is how you leverage AI as an integral part of your own data, workflows and solutions. It's important to leverage data and AI integrators, such as Eviden, so you can be ready for all opportunities across the information systems stack, and along the HPC (high performance computing)-to cloud-to-edge continuum.

6. **Stay agile and open.** Today's flourishing open-source communities are successfully competing and collaborating with traditional large independent software vendors (ISVs) in AI and GenAI. Traditional ISVs also increasingly propose open-source models. This not only offers users more opportunities to adapt the models to their own needs. It also dramatically accelerates ecosystem-wide technological progress by encouraging vibrant communities of contributors – again underlining the power of open innovation.
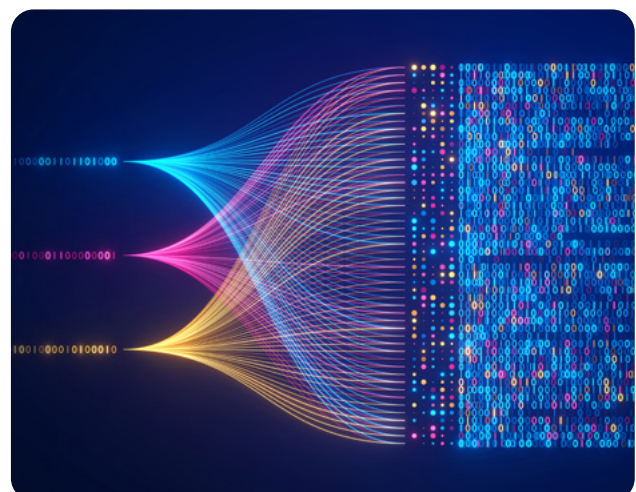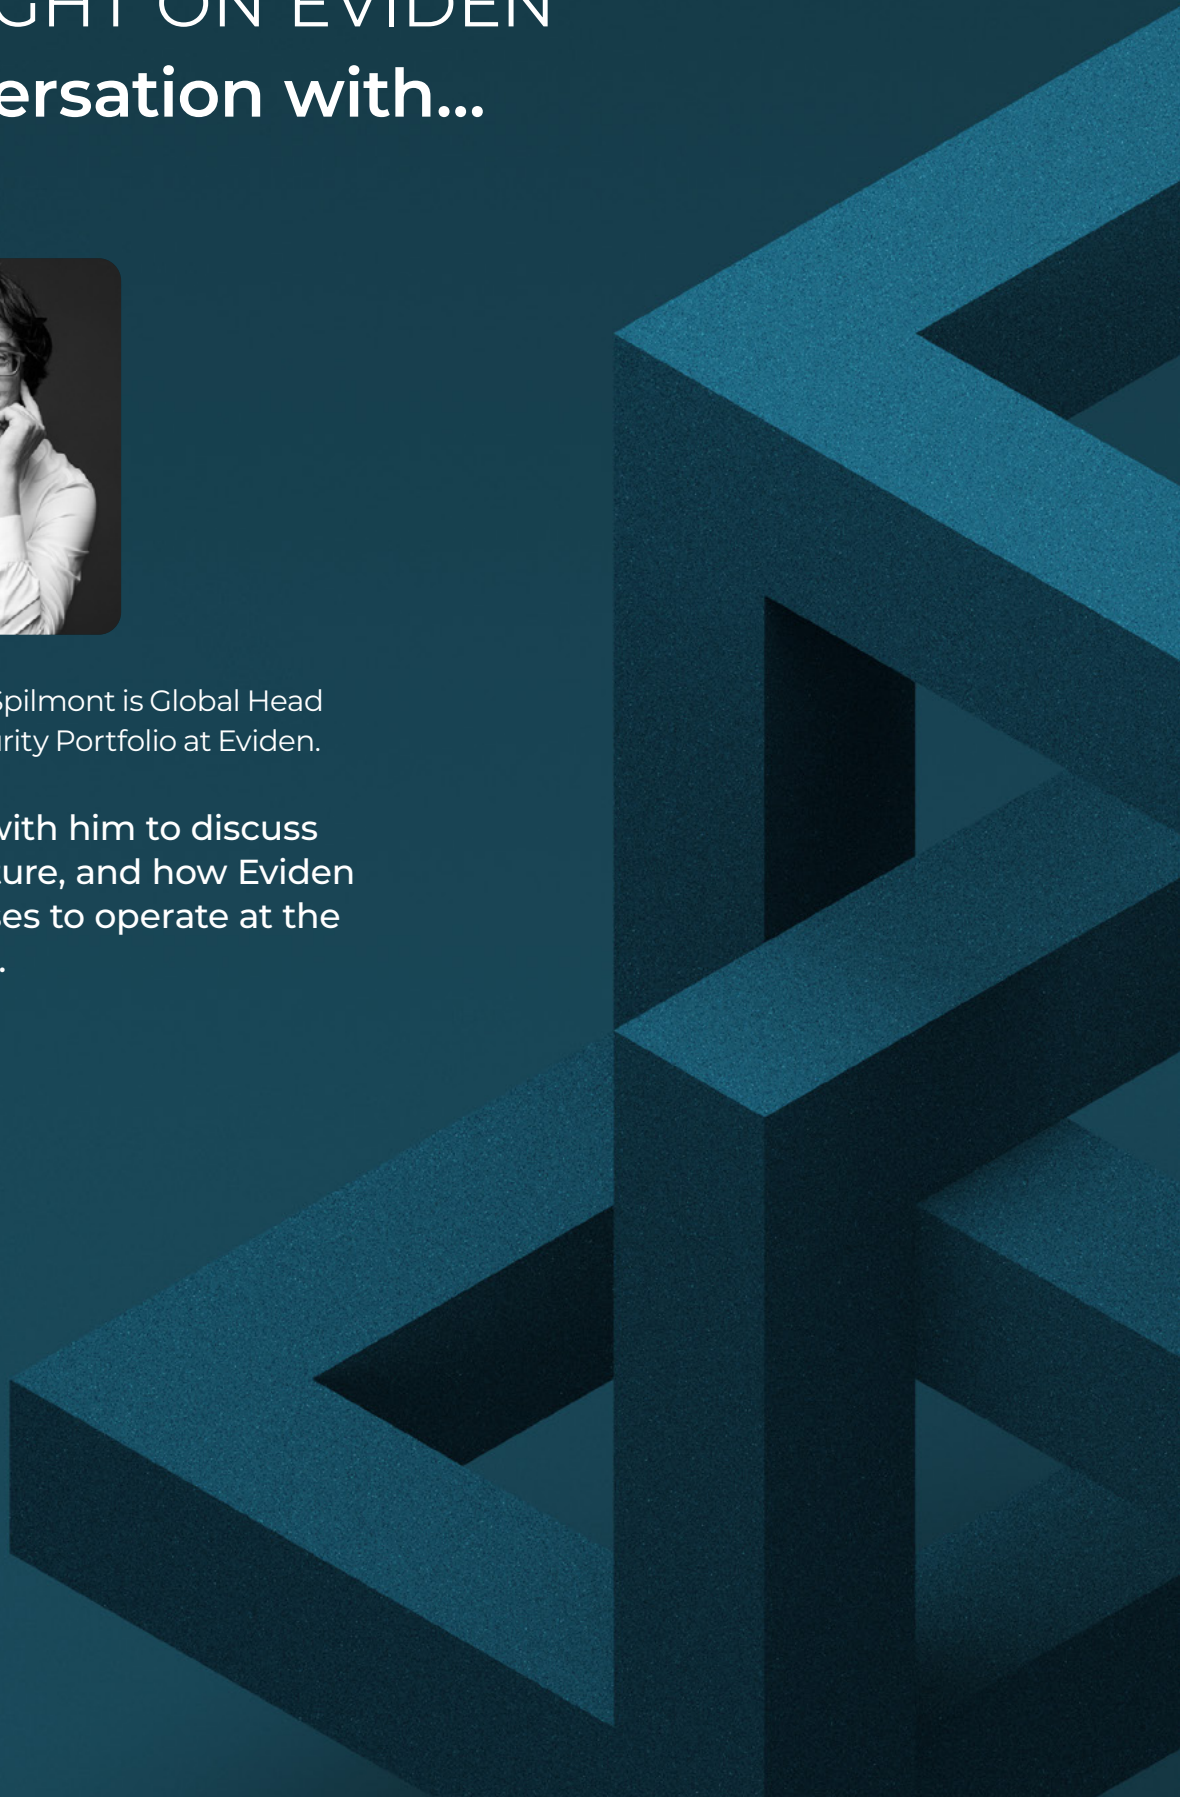
SPOTLIGHT ON EVIDEN
In conversation with...

Jean-Christophe Spilmont is Global Head of Big Data & Security Portfolio at Eviden.

We sat down with him to discuss AI today, its future, and how Eviden helps enterprises to operate at the new AI frontier.

What's Eviden's position on AI?

AI is in our DNA. Our heritage lies in data – we've been pioneering data platforms, edge/AI and high-performance AI for years. This new era is precisely why Eviden has been carved out from Atos's traditional business. We launched our Generative AI Acceleration Program to help businesses and organizations exploit, scale, and leverage the power of GenAI with complete trust. We've been recognized as a world leader by the first analyst survey of GenAI service vendors, conducted by HFS. That's why we position ourselves as a trusted partner for enterprises in their data/AI transformation.

Where do you see enterprises on this transformation journey? What challenges are they facing?

A number of analyst reports have found that over 70% of executives say they believe GenAI will disrupt their business in the coming years. Their challenges are pragmatic. How can they identify and deploy fast-to-value use cases right now? How do they avoid proof-of-concept purgatory by productizing and scaling their GenAI apps? How should they best tackle compliance, business sovereignty, and security risks?

Some just want to select a few ready-to-use vertical use cases. Others want to go deeper and integrate AI into their applications. Some want to adapt, and even build, GenAI models themselves or combine models to build specific applications. Others want to fine-tune performance with dedicated infrastructures, as-a-service or on-premise.

How does Eviden help?

We have developed a set of consulting programs and accelerators to meet these needs, with in-depth access to the best of the GenAI business and research ecosystem (AWS, Databricks, Google, Intel, Microsoft, Nvidia, Snowflake, international consortiums, and others).

For us, data (strategy, architecture, engineering, governance, sovereignty, responsibility) is the cornerstone of AI success. Then, we help customers leverage AI and GenAI in their information systems.

In a nutshell, we help customers transform their core business applications into powerful, trusted systems of prediction and execution.

How do you see the future of AI?

Among the key trends, we expect the explosion of open source GenAI models and specialized as-a-service vendors, the rise of on-premise dedicated infrastructures for training and inference, the evolution of Responsible AI, and the development of GenAI marketplaces. But for us, the key trend is the explosion of specialized AI models. Large applications such as ChatGPT or Gemini are often seen as monolithic. Yet they are often composed of multiple models in the background. For us, the future lies in the combination of multiple components, fine-tuned and orchestrated to meet specific needs.

What will this look like in practice?

A great example is the rise of autonomous agents. Initially dedicated agents will be deployed by people to complete tasks. But the future lies in agents autonomously collaborating among themselves, eventually giving way to what some analysts call 'autonomous business'.

In previous years, employees were augmented with robotic process automation (RPA). Today, they're working with Copilot. Tomorrow, swarms of autonomous agents will work for them. The social mobile, analytics and cloud (SMAC) era has disrupted customer relations. Tomorrow, customers will also have their own copilots. If you're a B2C vendor, maybe it's the future incarnations of Siri, Google Assistant or Alexa with which you'll have to negotiate. And after that, maybe your customers will be machines.

What are the timeframes?

It's impossible to predict precisely. But at the turn of this century, it took Netflix eight years to reach 100 million users. By the end of 2022, for ChatGPT, it took two months. In 2023, for Thread (Twitter's (now X') competitor), it took just four days!

We're just at the start of this revolution. It's said that people overestimate what can be done in six months, but underestimate what can be done in five years. At Eviden, we're here to help unlock the possibilities of AI for the benefit of people, businesses, society and the planet for decades to come.

[2]HFS Horizons Report, Generative Enterprise™ Services 2023

What's next?

Want to keep reading?

**Ignite your AI
journey Webinar**

**Your next story
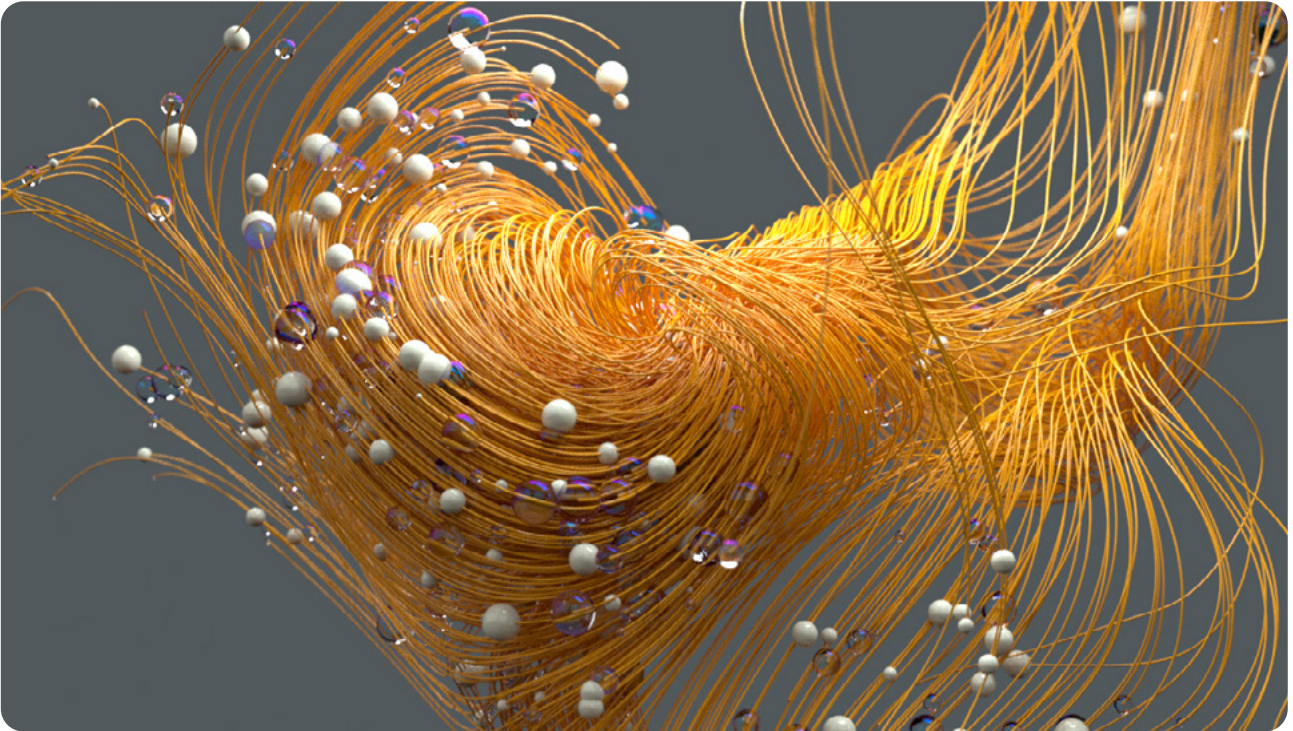unfolds with
Generative AI**

**Possibilities
magazine**

Want to talk to an expert?

Contact-us

Lexicon

Terms you may have heard (and some that may be new).

AI inference: the process of running live data through a previously trained AI model to solve a problem, take a decision or make prediction.

Deep learning: a form of AI whereby machines can process data in a way that's inspired by the human brain. Artificial neural networks use multiple layers of processing (hence deep) in order to recognize complex patterns in text, images, sounds and other data to produce progressively high levels of output.

Foundation (or base) model: a general-purpose machine learning model that's pretrained to carry out a range of tasks instead of bespoke applications or one-off use cases. This means that it provides a basis of knowledge that can be adapted and fine-tuned for specific requirements, making it an important step forward in enterprise AI.

Hallucination: AI hallucinations occur when AI learning models produce outputs that are completely erroneous or make no sense to human beings. These outputs are confidently presented by AI as facts, and are caused by problems within the training data or with how the learning model is working.

Large language model: a machine learning model trained on massive volumes of data in order to produce language and responses that sound human. This has transformed natural language processing by machines (which has existed for decades) into the GenAI applications we see today.
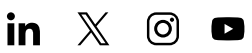
Liquid neural network: this recent innovation is a form of AI based on algorithms that can continuously adapt to new and changing streams of data. It has the potential to revolutionize AI applications that operate in a context of constantly changing or moving environments, such as drones, advanced robots and driverless vehicles. This recent innovation could also enable time-based decision-making based on data that is indexed in chronological order.

Machine learning model: the term to describe what a machine learning program learns from data. The more data that is used, the better and more accurate the machine learning model becomes.

Quantum AI: quantum computing (itself still emerging) offers vast improvements over traditional computing in computation time, quality, cost and energy usage. Based on quantum physics, quantum computing is non-binary – which means that unlike traditional computing, it doesn't provide a definite answer. Instead, it averages the right answer over multiple identical iterations and works with multiple states at the same time. Combine that potency with AI capabilities and it's hard to overstate the potential this will create. Still only nascent, early applications include financial trading.

Transformer: the tool that accelerated GenAI by making large language models exponentially more powerful. Transformers dramatically enhance and accelerate how machines can understand text. They can analyze whole sequences of language (paragraphs and documents), not just individual words. This makes it possible to capture context and meaning in a way that's faster, better and more efficient.

Connect with us