Digital Services

Generative AI

Visa for a more interactive digital future

orange™ Business

Introduction

It's hard to ignore the ChatGPT trend, as the tool has made headlines in the media since late 2022. Beyond the numerous associated use cases, it's the accessibility and the quality of its natural language that make OpenAI's solution stand out. While Generative AI is not new - and certainly not artificial intelligence - until now, it has only been the domain of specific experts and research laboratories. The real revolution in ChatGPT lies not so much in the technology itself but in its ability to target a broad audience.

Like social networking and other digital applications, Generative AI has quickly made its way to the walls of the enterprise. 63% of French companies are considering the introduction of Generative AI into their working methods.[1] The reason is that while only 13% of employees use Generative AI daily, almost half (45%) use it several times a week. Yet few people know precisely what Generative AI is or how it works. While 71% of French people have heard of it, only 36% understand it. As a result, most users are not necessarily aware of the risks of Generative AI in general and ChatGPT in particular. This lack of awareness can have dramatic consequences for the company: 68% of employees use ChatGPT at work without telling their employer.[2]

And the trend is not about to reverse. OpenAI, Microsoft, Google and Meta today will be joined tomorrow by Amazon and even Apple, which announced on July 20, 2023, that it had started work on Apple GPT, its chatbot service. Generative AI is set to redefine human-machine dialogue, disrupt the place and role of technology in the digital world, and, more generally, profoundly reinvent language production and creative processes. Companies no longer have a choice! Now is the time to get to grips with these new tools and regulate their use to limit abuses while facilitating adoption.

We are at the beginning of a new era in digital transformation. Generative AI is the forerunner of an unprecedented wave of innovation, which will question many business models and force companies to review their strategies.

That's why supporting its deployment and facilitating its adoption is vital.

This is why Orange Business is investing heavily in this technology, including training, development, partnerships, and prototyping. As a digital leader, the Group intends to be one of the pioneers in the market and a driving force in getting businesses to embrace Generative AI. It is for this reason that we have produced this white paper.

Aliette Mousnier-Lompré,
Orange Business CEO

1 According to a survey conducted by GetApp, a Gartner subsidiary, July 2023.
2 Fishbowl survey, February 2023

Contents

Development of Generative AI: Impact on businesses

In January 2023, two months after its launch, the ChatGPT website was attracting an average of 13 million unique visitors per day, with daily traffic growth of 3.4%.[3] The site welcomed one billion monthly visitors in April, including almost 100 million active users. Why are Generative AI and LLMs (Large Language Models) so popular? What impact will this have on businesses? Is Generative AI a simple evolution or a real revolution? Let's find out.

"ChatGPT and other major language models are the result of decades of inputs from various people."

To better understand how Generative AI will transform business practices, it is essential to understand what it is all about. The term "Generative AI" suggests that this AI will generate content. However, this is only partially true, as it works in both directions: LLMs are undoubtedly capable of generating content, but they are also capable of "digesting" content. This 2nd aspect is crucial because it sheds light on how they work and the risks they present for businesses.
At the same time, we need to distinguish between different types of Generative AI. Thanks to ChatGPT, the best known is textual Generative AI, which analyses and produces text. Based on the same model, other Generative AIs will work in the visual, musical, or audio domain. In this section, we will focus on the most mature of these Generative AIs, namely textual AI, or LLMs, whose impact on the workplace is currently the most significant and impressive.

From a simple technological innovation...

From a technological point of view, LLMs are more an evolution than a revolution, apart from the fact that they combine up to four learning modes[4] in the same artificial intelligence:

- Unsupervised mode, or embedding, consists of assigning to each word a vector (a series of numbers) that best defines this word in a mathematical space called a vector space.
- Self-supervised mode uses part of a sentence to complete the sentence with what makes the most sense from a statistical point of view.

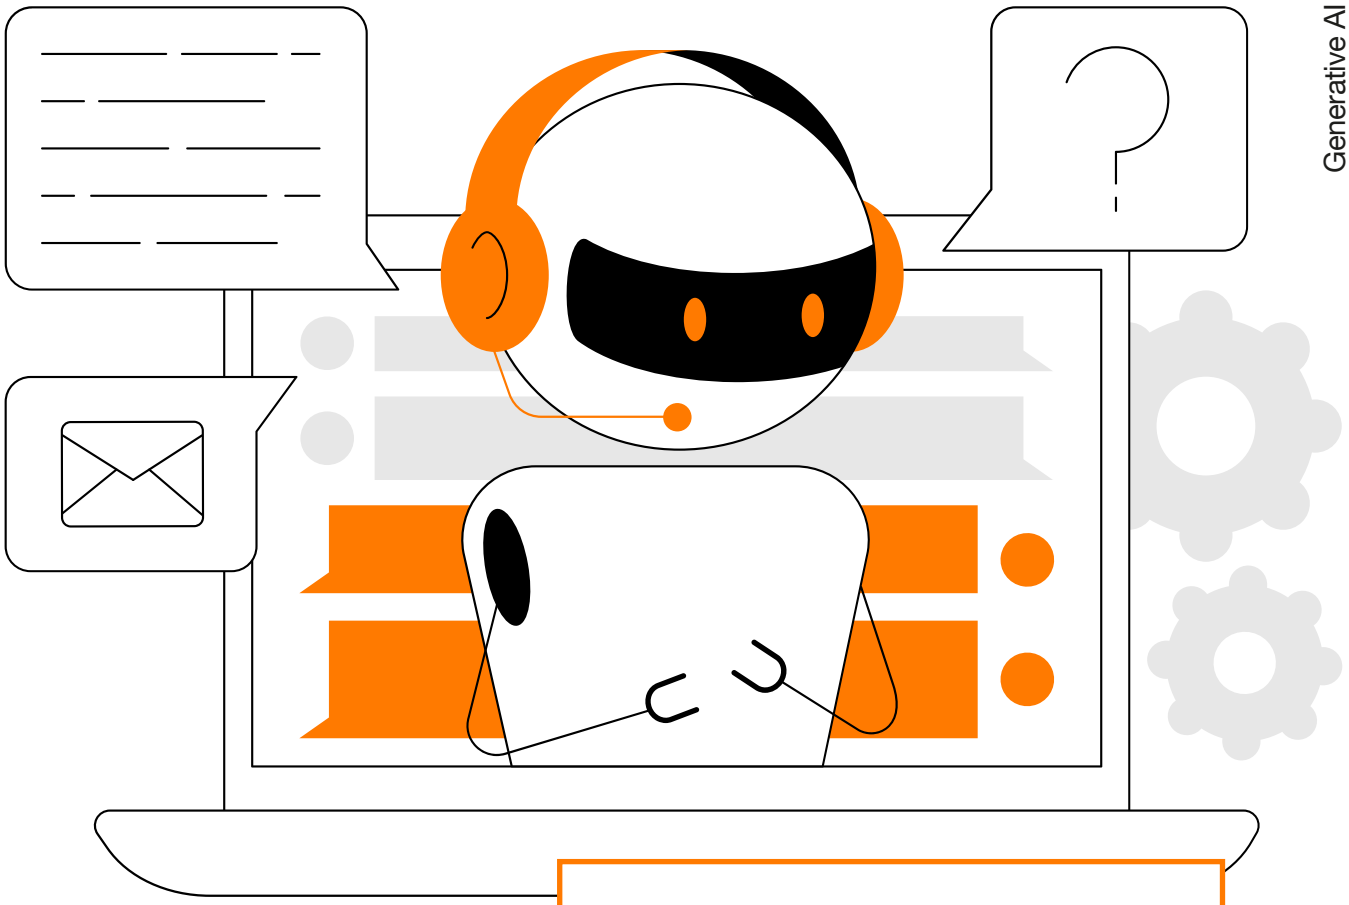

- Supervised mode (equivalent to the InstructGPT layer in OpenAI) teaches the model how to answer questions.
- Reinforced mode to rank the answers in order of preference and evaluate the most appropriate. This stage requires human intervention.
- There is also an additional security layer to identify prohibited answers and questions.

Two other technical developments are also significant. The first is based on transformer technology, which, ironically, was invented in 2017 by... Google. In its article entitled “Attention is all you need”[5], Google Research explained how to build a transformer to train deep learning models with self-supervised learning.
The next innovation is closely linked to the 2nd: it is the ability to run transformers on a computing infrastructure with a gigantic corpus of training, the heart of this infrastructure being based on GPUs (Graphical Processing Units), notably of the H100 type manufactured by Nvidia.

Generative AI: evolution or revolution?

Contrary to what you might think, Generative AI is not so recent. It began to emerge with the rise of deep neural networks. In 2014, generative adversarial networks (GANs) could generate images resembling real photographs. Generative AI is the direct result of decades of research into artificial intelligence. As Yann LeCun, Meta’s head of AI, points out, “ChatGPT and other major language models did not come out of nowhere. They are the result of decades of inputs from various people.”
But the catalyst was the release of ChatGPT to mass-market users by OpenAI in November 2022. Why? Because it raised awareness about the impacts generated by using such a tool.

“Generative AI pushes back the limits of what machines can learn and create.”

… to a usage revolution

Although the main innovations in Generative AI are based on combining these four learning modes and transformers, the real revolution lies in how it is used. LLMs mark the start of a new era in artificial intelligence and man-machine dialogue, promising a revolution as big as the Internet. From now on, anyone can interact with a computer using natural language. In response, the machine can interpret and execute the user’s requests. Generative AI is, therefore, capable of creating new data from a corpus of training data, a prompt (user request) and a context.
For example, GPT-4 can give advice in the text field, draw up a plan, write articles, product sheets and even code.

Unmatched computing power

So, if Generative AI isn’t technically innovative, why is it getting so much hype? Because its unequalled language processing power opens infinite possibilities for users and, consequently,

Expert opinion

Didier Gaultier, DataScience & Ethics AI Director, Digital Services France, Orange Business

Vapnik's law

Vladimir Vapnik is a Russian mathematician and computer scientist known for developing support vector machines (SVMs) and originating the Vapnik-Chervonenkis statistical learning theory. What does this law say? You need a good balance between the model's complexity and the training data's complexity.
In short, if we run out of data, it will not be possible to increase the complexity of the model and, ultimately, improve it or extend its capabilities. As the amount of text available on the Internet is limited, it is impossible to increase the intelligence of models indefinitely. In this sense, Generative AI (on its current architectures) will soon reach its limits.

for businesses. Generative AI pushes back the limits of what machines can learn and create: new product designs, personalised marketing content, sales pitches incorporating a wealth of information, assistance with medical diagnosis or forensic analysis, etc.

How Generative AI will change the game

It is difficult to predict all the new professional uses that Generative AI will promote or improve, as the field of possibilities is so broad. Conversational AI, on the other hand, must remain no more than a tool, as the machine lacks intention and accurate understanding; for example, it is not capable of distinguishing between what is ethical and what is not.
In some general uses, we should also consider that the best LLMs have error rates that can exceed 30%, which can be significantly reduced by specialising the scope of training or the context in which the LLMs are used. Artificial intelligence does not replace humans; it simply takes over specific tasks, increasing efficiency.
However, there are at least two areas that Generative AI is likely to disrupt significantly. Firstly, marketing, by enabling large-scale personalisation. Secondly, search engines, by being able to provide accurate, documented answers, force players like Google to rethink their business model.

Expert opinion

Michael Deheneffe, Strategy, Marketing & Innovation Director, Digital Services France, Orange Business

"Generative AI represents a major technological breakthrough. Simply put, computer science has known how to make calculations for 50 years. Today, it can write, learn, and synthesise. Of course, this does not mean thinking or understanding, but the boundary is becoming increasingly blurred. This is the first time there has been such a rapid disruption.
Another finding is the massive technology adoption by various profiles, even within companies. In comparison, while the quantum computer represents a real technological breakthrough, its adoption is nothing like it. The same goes for Metaverse.
Finally, there is the question of accessibility. Generative AI requires vast calculations, which are carried out in the cloud. So, let's take the three terms of the equation, i.e., major technological breakthrough, hyper-rapid adoption, and widespread accessibility. The emergence of Generative AI can be seen as akin to the internet revolution. And as with any major technological breakthrough, its short-term effects are overestimated, and its long-term effects are consistently underestimated.
The trend is sure to last longer than predicted. As a result, companies are caught between two central tendencies. On the one hand, they want to test the technology. Still, on the other, they want to do so securely and guided, fearing issues of confidentiality, intellectual property, ethics and so on."

The impact of Generative AI on businesses

Companies adopting Generative AI could benefit from improved productivity and reduced costs, two major competitive factors. The outcome is the ability to generate new content in response to competition rapidly. One thing is sure: Generative AI will transform how we create and consume content. In this sense, it can represent a strategic differentiator for businesses. On the other hand, Generative AI also raises several questions and challenges. For example, what about the copyright on the works used to train it? It is difficult to know precisely which data certain LLMs have been taught, raising intellectual property issues for companies that use them.

The limits of Generative AI development

Some limitations are already emerging, the first of which is methodological. This is because models such as GPT-4 or its successors need a large proportion of the content already published online for their training. Increasing their power and complexity would give them access to even more content (see box on Vapnik's law), limiting their development potential. In addition, many resources that would be valuable for training LLMs are not published on the web and, therefore, remain inaccessible to Generative AI. As a result, general LLMs alone cannot provide exhaustive and relevant answers in most specialist areas. This long-term process will involve specialising the training of certain LLMs on specialised corpora and digitising more and more data.
Another hurdle is creating content that is always up to date using Generative AI. We know, for example, that ChatGPT does not currently include data from after September 2021.[6] The plugins can certainly fetch content from the Internet, but they integrate the data collected either in the prompt or in the context, but not in the training corpus.

“A number of limitations are already apparent, the first being a methodological one.”

As a result, the tool’s ability to analyse the most recent data is also limited. In this sense, the updating of training data represents a fundamental limitation. And that’s not counting the biases and other fake news included in the training data, potentially penalising the answers’ relevance.
Like all technologies, the responsible and ethical use of Generative AI is essential to maximise its benefits while minimising the risks. This is why businesses should opt for secure environments dedicated to them, such as Azure OpenAI (Microsoft) or ChatGPT Enterprise (Open AI), and avoid a mass-market product like ChatGPT, which does not offer the necessary level of confidentiality.

3 ChatGPT Statistics 2023, ToolTester, July 2023 https://www.tooltester.com/en/blog/chatgpt-statistics/
4 For more details on these learning methods, see our white paper “Artificial Intelligence - Stay in control of your future”, March 2022
5 Attention is all you need”, Google, 2017 https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
6 This is data from 2023

The core technologies of Generative AI

Generative AI is not a recent concept. The neural networks on which it is based have been known for decades. So why did we have to wait until 2022 to hear about this technology? We'll give you a few explanations.
Before going into the details of the technologies, it is essential to note that we will focus on the core technologies of ChatGPT, which, in 2023, remains the most mature fundamental model of Generative AI in terms of processing capacity, generation capacity and creativity.

First wave: AlexNet & deep learning

It all began in 2012 when Alex Krizhevsky implemented neural networks on graphics cards. Why was this revolutionary? Neural networks are highly greedy in terms of computing resources and data volume. They are very easy to "overfit", i.e., to over-interpret.[7] Until then, the only way to avoid this was to reduce the number of settings and, ultimately, the model's performance. Graphics cards have made it possible to cut calculation times by a factor of 100. At the same time, huge datasets, such as ImageNet, were made available, helping to improve

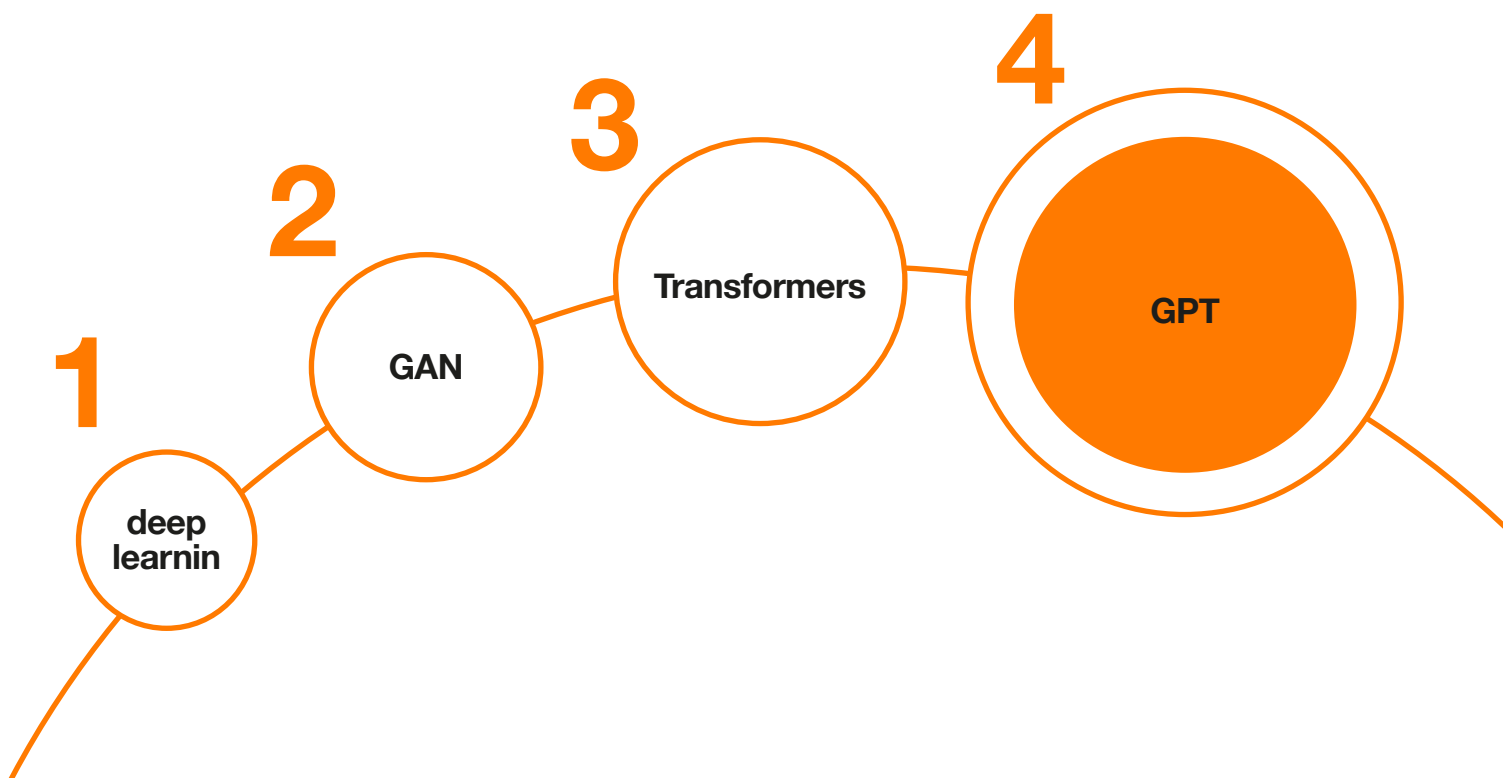

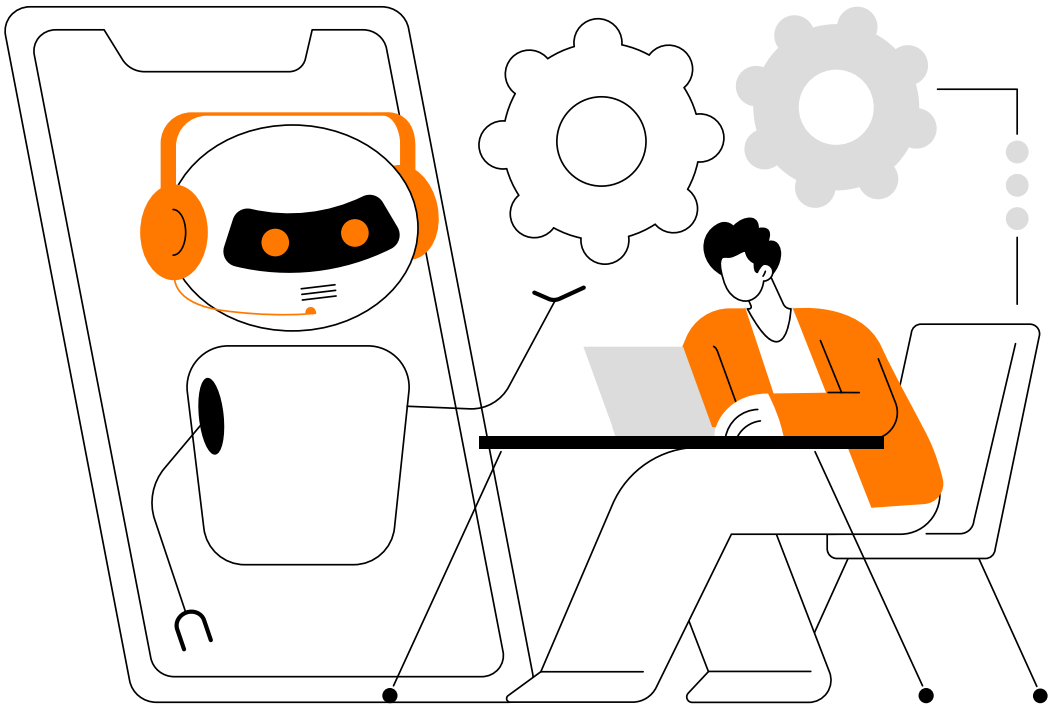

model performance. These performances were made public for the first time at the 2012 ImageNet competition. Often cited as a reference in computer vision, this architecture, called AlexNet, is considered the first convincing approach to deep learning, bringing to life decades of research sometimes carried out against all odds during the “AI winter”. This step made it possible to release a large amount of research in applied mathematics, computer science, neuroscience, etc., which has encouraged the development of Generative AI.

Second wave: GANs

In 2015, the arrival of GANs (Generative Adversarial Networks), thanks to Ian Goodfellow, made it possible to put a set of architectures and methods in place. The principle consists of pitting two neural networks against each other: one that generates unsupervised images and another that tries to determine whether the image is valid. This stage can be considered the first high-performance generation of Generative AI. However, GANs resulted in particularly unstable architectures.

Third wave: the transformers

Launched in 2017, the transformers models developed by Google Brain have the distinctive feature of being based on a concept known as attention. The model consists of implementing a sequential selection of information as processing progresses. These transformers began to be used in natural language processing (NLP) with good results, particularly in the case of BERT-type models, which are still used in many applications today.

Fourth wave: GPT 3 and 4

Since then, significant groups such as Microsoft (OpenAI), Meta (Facebook) and Google have decided to continue developing these models to work on Generative AI. By adding more and more neurons and parameters, and above all, by combining various learning modes, OpenAI has succeeded in creating its ChatGPT solution. Its performance is based, among other things, on combining the four learning modes described below and using Nvidia H100 graphics cards.

“ChatGPT remains the most mature fundamental model of Generative AI in 2023.”

Self-supervised mode

In 2018, the GPT model was developed, using self-supervised learning on text with transformers. This approach requires much less supervision of the data but involves text embedding and a large amount of unlabelled data. It also includes language learning from a particular corpus of text based on massive use of what has been published online. This stage explains why ChatGPT can operate in several languages. It should be noted that this stage is also the most computationally intensive, its sole aim being to complete the next word in a sentence or to fill in a missing word in a sentence.

Expert opinion

Michael Deheneffe, Strategy, Marketing & Innovation Director, Digital Services France, Orange Business

"Today, we can already say that Generative AI will be able to synthesise large volumes of unstructured data. In this sense, the technology can be seen as a junior personal assistant, capable of carrying out certain tasks but without the level of seniority needed to validate the level of detail and accuracy of what it advances. Every employee equipped with a Generative AI tool will benefit from invaluable help daily when it comes to, for example, a particularly time-consuming task. Beyond the technological aspect, it is in how we approach projects and the relationship between the company and its employees that Generative AI is revolutionary, transforming the trust between an employee and his employer or between a service provider and its customer. As a result, trust and authenticity will have to be considered to a greater extent when 'certifying' documents."

The supervised layer

More recently, in 2022, InstructGPT was introduced. This is a GPT model trained on conversations in supervised mode. This approach teaches GPT to answer questions. It involves providing the model with a certain number of questions and the expected answers. The complexity then lies in the context the AI needs to answer the question correctly. The context contains the instructions we give the model to enable it to respond. It is hidden in ChatGPT and included in the InstructGPT layer, which has undergone supervised learning. The result is a chatbot-type artificial intelligence, trained on a vast corpus of text and capable of answering any question in the raw, as long as there are elements of the answer in the training corpus.

Enhanced learning

While the first three learning modes create the "substance" of the response, this mode promotes the "form" and enables the AI to respond in a more relevant way. The supervised layer results in a somewhat rustic chatbot devoid of nuance. Here, a team of people will ask ChatGPT many questions, ask it to generate several different answers per question, then rank them in order of preference and even eliminate inappropriate answers in the form of reinforced learning with RLHF (Reinforcement Learning from Human Feedback). While this stage is not technically complex, it does require significant human resources.

Fundamental models in GAFAM hands

By increasing the size of the datasets and avoiding the need for massive data labelling, it was possible to create the first foundation models, which then, through successive refinements, made it possible to develop interactivity close to human interaction thanks to the InstructGPT layer. Today, the major groups such as OpenAI with Microsoft, Google or Meta are the only ones with the computing power needed to train multilingual fundamental

“In the future, there may be a lack of textual training data for the models.”

models based on increasingly gigantic transformers. This is why future developments in the basic Generative AI models partly rely on the GAFAMs.
However, it is essential to remember that a model that includes many parameters but lacks the supervised layer will appear relatively unusable compared to a simpler model that benefits from supervised learning.

Generative AI will soon reach its limits.

In September 2022, DeepMind (a paper published by Chinchilla) put forward new data scaling laws (also known as Chinchilla or Hoffman scaling laws) for optimal LLM (Large Language Models) data[8]:

- A corpus of 1.4T tokens[9] would be needed to optimally train an LLM incorporating 70 billion parameters, i.e., around 20 text tokens per parameter

Chinchilla says future Open AI models may require more than 10 times as much training data as their precursors and similar models. For information, it is estimated that GPT 3 already contained 175 billion parameters[10]. As a result, we run the risk of running out of textual training data for the number of parameters we want to add to a model.
So, there is no point in increasing complexity exponentially to improve model performance significantly. The future of artificial intelligence lies elsewhere, in architecture. Why not consider discovering a more efficient architecture than transformers in the future? For example, progress in fundamental mathematics (particularly high-dimensional vector spaces) could perhaps lead to further advances in the architectures currently used in AI.

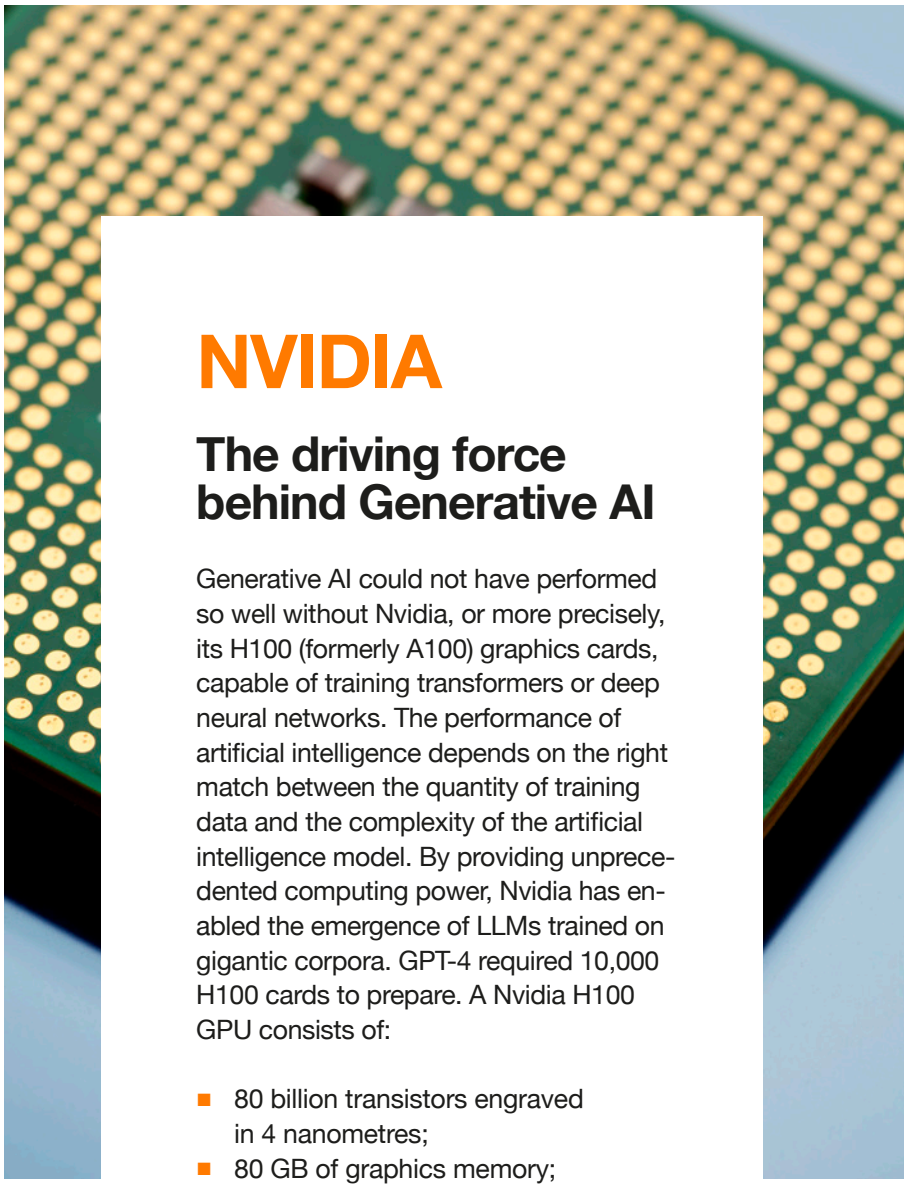

NVIDIA

The driving force behind Generative AI

Generative AI could not have performed so well without Nvidia, or more precisely, its H100 (formerly A100) graphics cards, capable of training transformers or deep neural networks. The performance of artificial intelligence depends on the right match between the quantity of training data and the complexity of the artificial intelligence model. By providing unprecedented computing power, Nvidia has enabled the emergence of LLMs trained on gigantic corpora. GPT-4 required 10,000 H100 cards to prepare. A Nvidia H100 GPU consists of:

- 80 billion transistors engraved in 4 nanometres;
- 80 GB of graphics memory;
- A bandwidth of several terabytes per second.

Finally, while machine learning is not intended to reproduce how the human brain works faithfully, some mechanisms studied by neuroscientists can provide new ideas for future AI research, such as the first neuron models, hierarchical models of vision and attention mechanisms.

7 To better understand the principle of over-interpretation, read our white paper “Artificial Intelligence - Stay in control of your future”. of your future” March 2021
8 Chinchilla data-optimal scaling laws, February 2023 https://lifearchitect.ai/chinchilla/
9 Word ≈ 1 to 1.5 tokens
10 GPT-3 on Wikipedia https://en.wikipedia.org/wiki/GPT-3

The leading players in the Generative AI market

The Generative AI market is still very new. As a result, the number of players will still be limited in 2023. Nevertheless, two main categories stand out. GAFAMs on the one hand, and startups on the other. Why are there so few mature players in this field? Because the heart and soul of artificial intelligence lies in access to data.

This is a critical point in Generative AI: to develop and train LLM-type models, the most important thing is the volume and quality of the data available. But even giants like Apple and Amazon face a relative lack of data, as they do not directly control an open online ecosystem on a global scale like Microsoft, Meta, or Google. As a result, a small number of companies currently share the market's potential. It is impossible to take an exhaustive look at all the players. We have therefore selected the following solutions, which are among the most advanced in LLMs and Generative AI.

Startups and Generative AI

Of course, startups are not the best equipped with data. They are still in the investigation phase, but their strong capacity for innovation and their native agility mean they can make rapid progress on the subject. In this section, we have therefore chosen to focus on four "promising" startups: Mistral AI, LightOn, Aleph Alpha & Hugging Face.

Mistral AI
Co-founded by Arthur Mensch (CEO) and Guillaume Lample (Chief Scientist and former Research Scientist at Facebook under Yann Le Cun), Mistral AI is a French company. It aims to create and share cutting-edge AI models capable of generating high-quality natural language from any input (text, images, audio, etc.). In July 2023, Mistral AI successfully raised significant funds.

LightOn
Florent Krzakala, Igor Carron, Laurent Daudet and Sylvain Gigan co-founded this other French company. LightOn provides essential augmented work and creativity models by leveraging exa-scale hardware and deep learning insights. Their flagship product, Para-

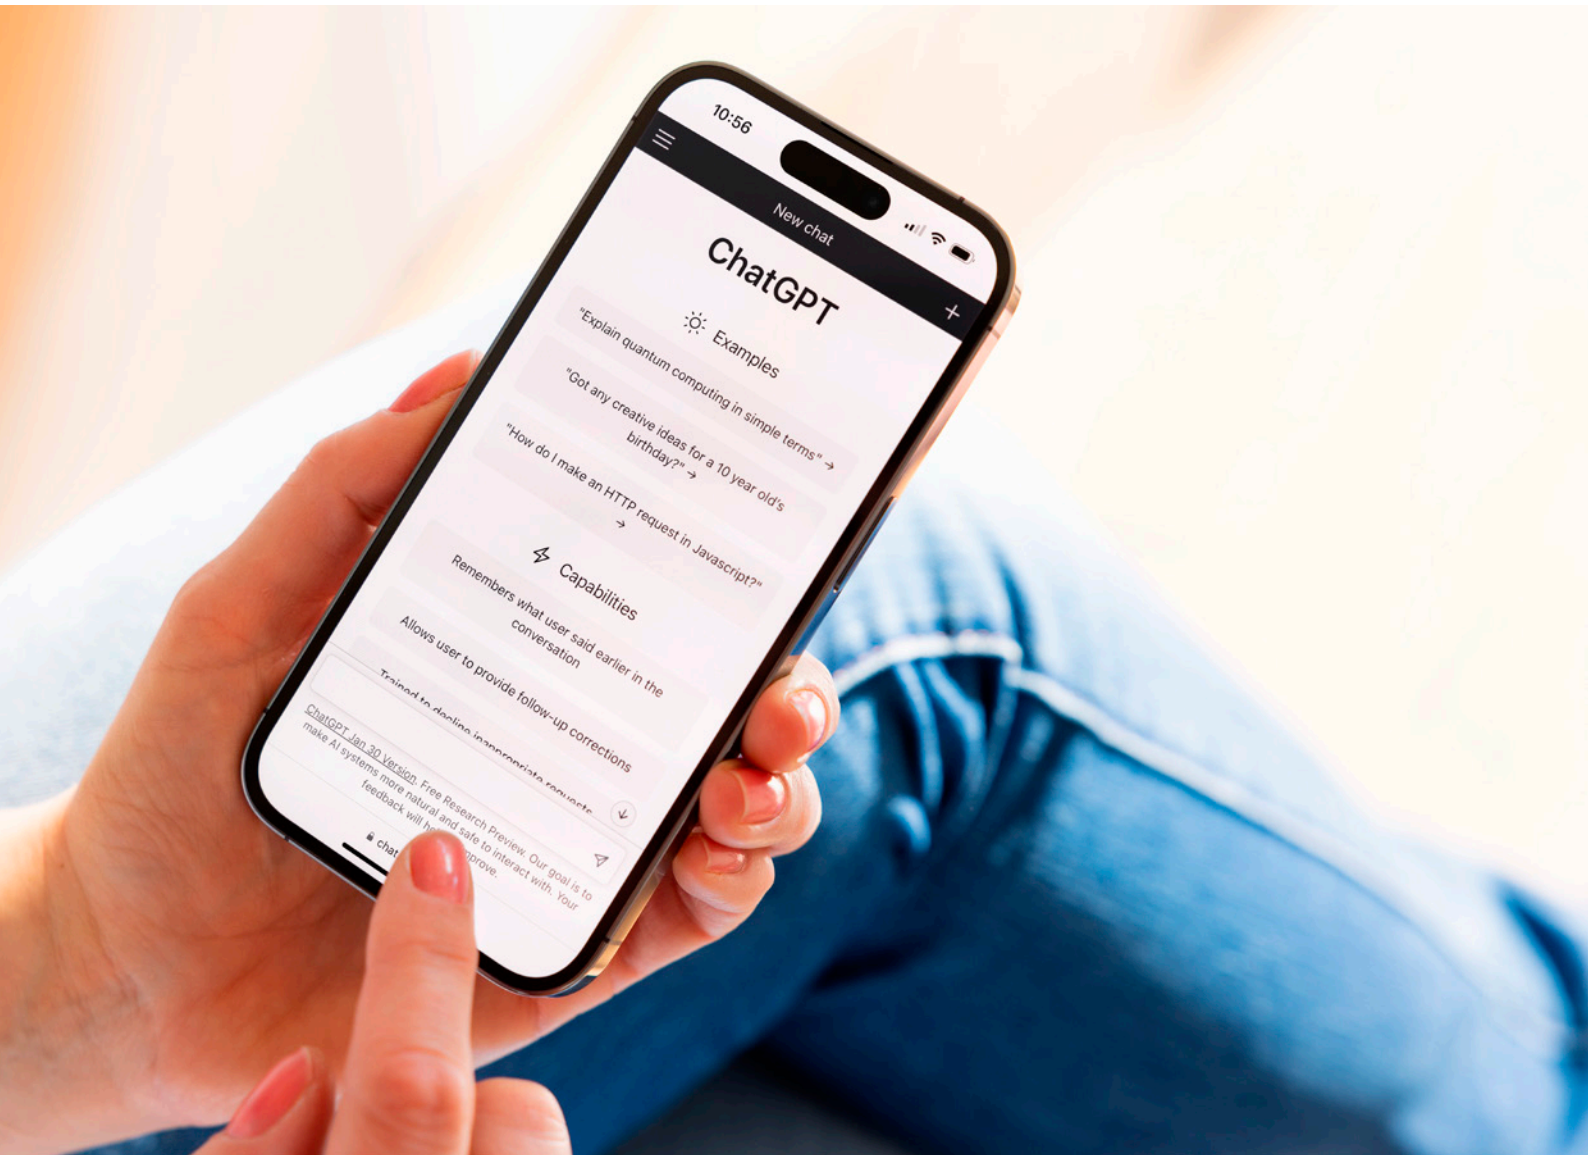

digm, is a robust and reliable solution for organisations using hosted language models for Natural Language Processing (NLP) tasks. They offer the benefits of private models, such as data control and adapting the model to specific needs. One of the many benefits of their solution is that it opens the door to transfer learning to private and sovereign data corpora. LightOn has partnered with Aleia, another French startup, to develop a secure European industrial offering of large-scale language models. Available at the start of the 2023 academic year, this offering aims to revolutionise sovereign industrial artificial intelligence by integrating secure, easily deployable workspace and high-performance large-scale language models.

This is a critical point in Generative AI: to develop and train LLM-type models, the most important thing is the volume and quality of the data available.

Aleph Alpha

Germany is also working on Generative AI in Europe, notably through the German startup Aleph Alpha, co-directed by Jonas Andrulis and Samuel Weinbach. With their Luminous solution, Aleph Alpha has reached the first milestone on the road to Generative AI that is powerful in content, explainable, sovereign, and trustworthy. Aleph Alpha continues to work towards

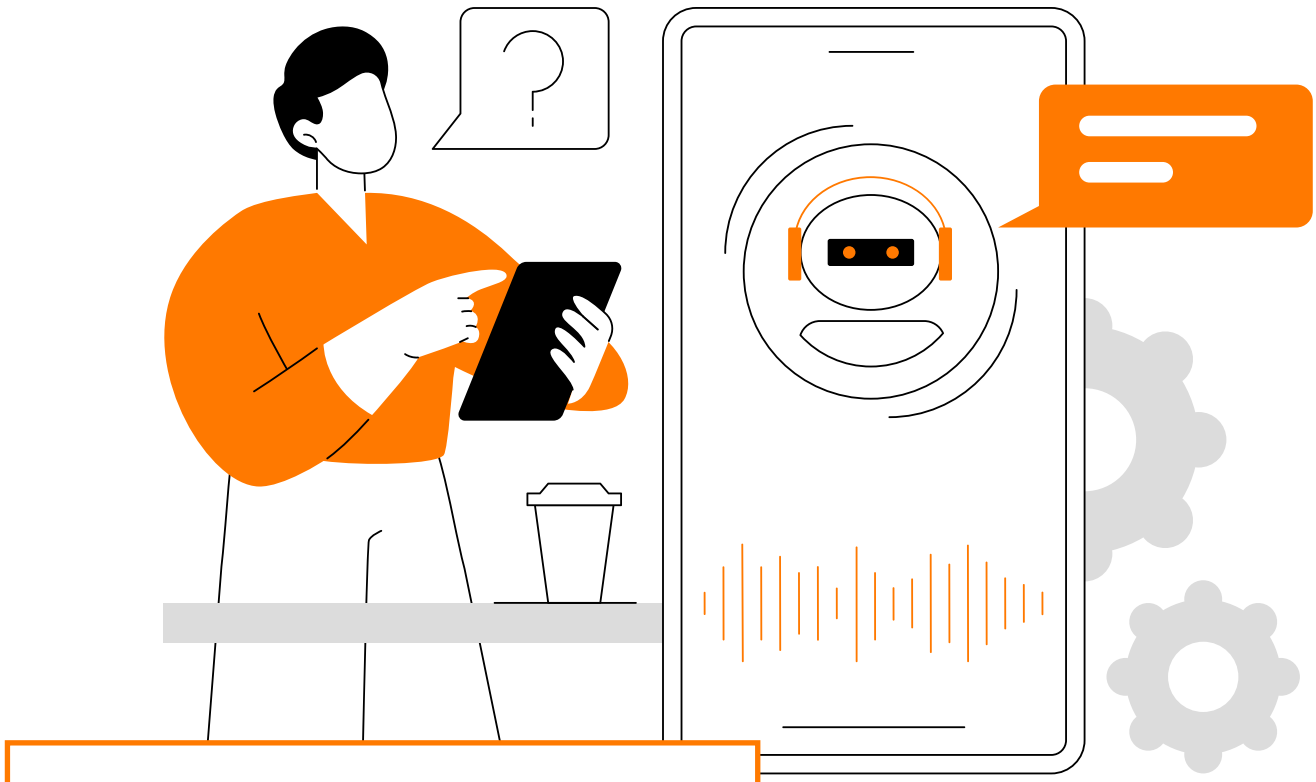

Bing chat

Bing, Microsoft's search engine, saw a 15% increase in traffic after the GPT integration, while the Bing mobile app was downloaded 750,000 times, with a peak of 150,000 daily downloads during the first week of integration in May 2023.

greater transparency in operating large language models. It should be noted that in 2023, Aleph Alpha is mainly focused on the German market and is working in France in partnership with HPE (Hewlett Packard Entreprise).

Hugging Face

Hugging Face is unique because it is an American company (its head office is in the United States) but was founded by Frenchmen Clément Delangue, Julien Chaumond and Thomas Wolf. Hugging Face is, first and foremost, a community and offers a data science platform that provides tools enabling users to build, train and deploy models based on open-source code and technologies. They are the creators of Transformers, an open-source library for building Generative AI models. Several Hugging Face models have already been integrated into the Microsoft Azure ecosystem.

GAFAMs at the forefront of Generative AI

It is impossible to talk about Generative AI without mentioning the hyperscalers, starting with OpenAI and Microsoft (GPT-4), Meta (LLaMA 2) and Google (PaLM 2).

OpenAI et Microsoft

By investing ten billion dollars in the American startup, Microsoft has played a significant role in making OpenAI the current leader in Generative AI. While ChatGPT is now integrated into the Bing search engine and Azure, Microsoft has nonetheless added its own layers of security to the consumer solution. Using ChatGPT implicitly means accepting that the information shared in the prompts can be re-used later to train the models and, therefore, exposed to other

customers. In this sense, the consumer offering can even re-use potentially personal or sensitive data that has been shared with it to deduce future responses. To this end, Open AI has launched the "ChatGPT Enterprise" offering, which now guarantees professional data confidentiality.
The most secure solution, however, is Azure OpenAI, which integrates the confidentiality of prompts and contexts and provides a guarantee that extends to all source documents. Microsoft thus guarantees its customers complete and contractual confidentiality when using Azure.

Google PaLM2 / Bard
Considered to be Microsoft's main competitor in Generative AI, Google has long held the top spot among the web giants thanks to the power of its search engine. But today, its business model could well be entirely disrupted by Microsoft.
In response, Google launched PaIM2 / Bard, an artificial intelligence text generator, at the beginning of 2023. The chatbot is now available in Europe.

Meta
In February 2023, Meta also entered the Generative AI race by making its language model, LLaMA (Large Language Model Meta AI), available to the research community. More recently, LLaMA2 was released in open-source mode for commercial use.
Based on smaller but powerful models, LLaMA2 could reshuffle the cards for LLMs, as Bing ChatBing, Microsoft's search engine, recorded a 15% increase in traffic after the integration of GPT. At the same time, the Bing mobile app was downloaded 750,000 times, with a peak of 150,000 daily installations during the first week of integration in May 2023. Many companies could, therefore, be tempted to train the model with their corpus of reference texts.
The advantage of this would be greater control over sources and greater transparency and control, particularly for chatbot applications. As a result, we are likely to see many LLaMA2 variants (or forks) flourish over the coming month.
With open-source solutions from Meta and Hugging Face, startups working on highly promising solutions, the advances made by the GAFAMs and their ability to challenge each other, the LLM market looks set to be very dynamic. In April 2023, AWS responded by presenting Amazon Bedrock in response to Azure OpenAI. With a service advertised as very affordable, AWS also ensures the security and confidentiality of business users' data. At the same time, Amazon has launched an acceleration programme for ten startups working on Generative AI. Even Apple, at its own pace, seems to have recently entered the Generative AI race with AppleGPT. As usual, when it comes to R&D, Apple is keeping a low profile because it's clear that, with SIRI getting a bit old, the competitive stakes in Generative AI will be high for the Apple brand in the years to come!

"By investing ten billion dollars in the American startup, Microsoft has played a major role in making OpenAI the current leader in Generative AI."

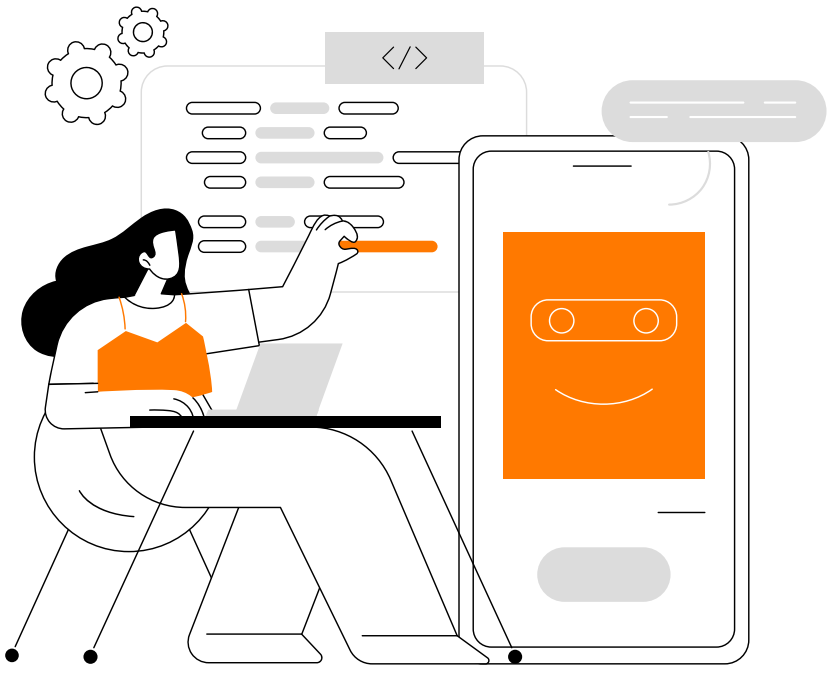

Impacts & challenges of Generative AI on businesses

49%

believe that Generative AI will radically transform their businesses.

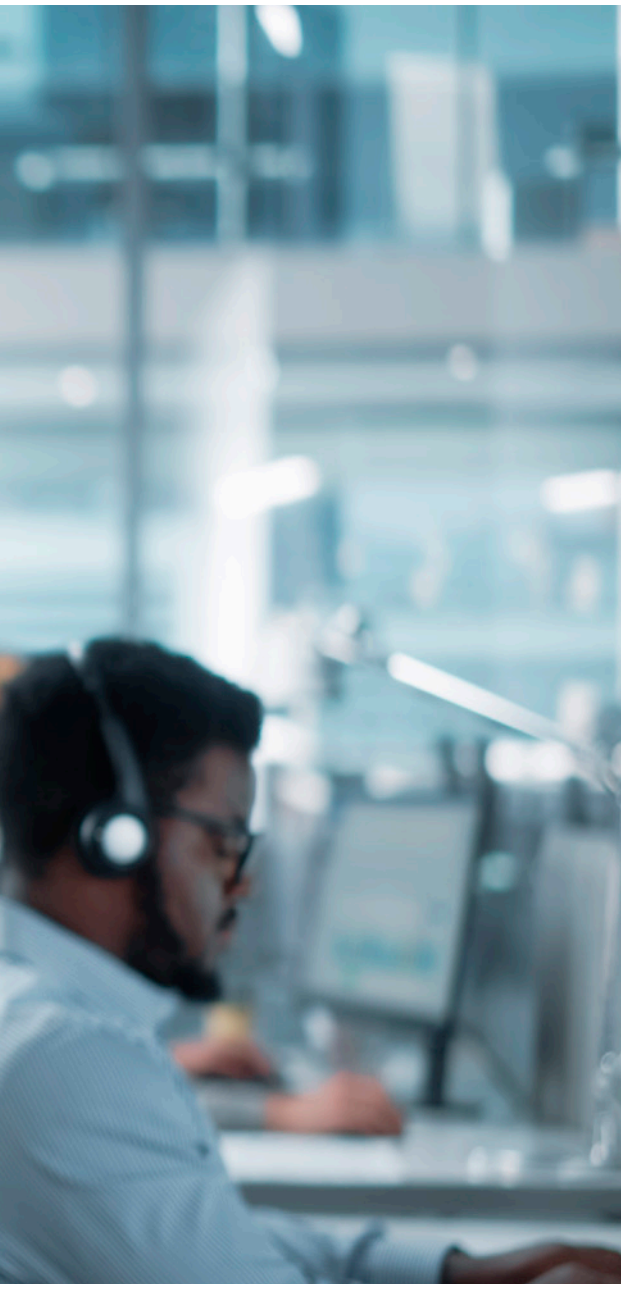

ChatGPT, Dall-E, Midjourney. For 52% of the French, Generative AI represents a new industrial revolution.[11] And 49% believe that it will radically transform their businesses. But what exactly are the challenges and impacts of this technology on the business world?
While there are many possible uses for Generative AI, it is mainly used to speed up several recurring tasks or to assist employees in their jobs by taking on repetitive or time-consuming tasks. To explain what we mean in this section, we will focus on a use case in human resources.

11 The French and Generative AI survey, Ifop for Talan, May 2023

What are the main challenges of Generative AI in business?

One of the first challenges of Generative AI in companies is to save time on low-added-value tasks. The aim is to enable teams to concentrate on their expertise and core business.

Reduce the number of recurring and administrative tasks

This is the case for IT development teams, for example, who can rely on Generative AI tools to generate simple code, compile code automatically or pre-fill a certain amount of information. The same is true for HR teams: Generative AI is capable, based on a certain number of documents, of proposing a chatbot to answer administrative questions relating, for example, to union agreements, holiday entitlements or employment contracts. As a result, HR teams can concentrate on the company's human relations (company life, transformation, etc.) and developers on algorithms and complex programming.

Quickly generate content

Content generation is a significant issue for marketing, communications, and sales departments. There are numerous practical applications: social media posts, product sales descriptions, presentations based on information from various documents, extraction of critical information, generation of general contractual clauses, technical code documentation, etc. Generative AI can also address specific branding needs (slogan, logo, visual or even video illustration) or offer the same presentation in different forms to differentiate from the competition. The result: productivity gains for teams and competitive advantages for the whole company.

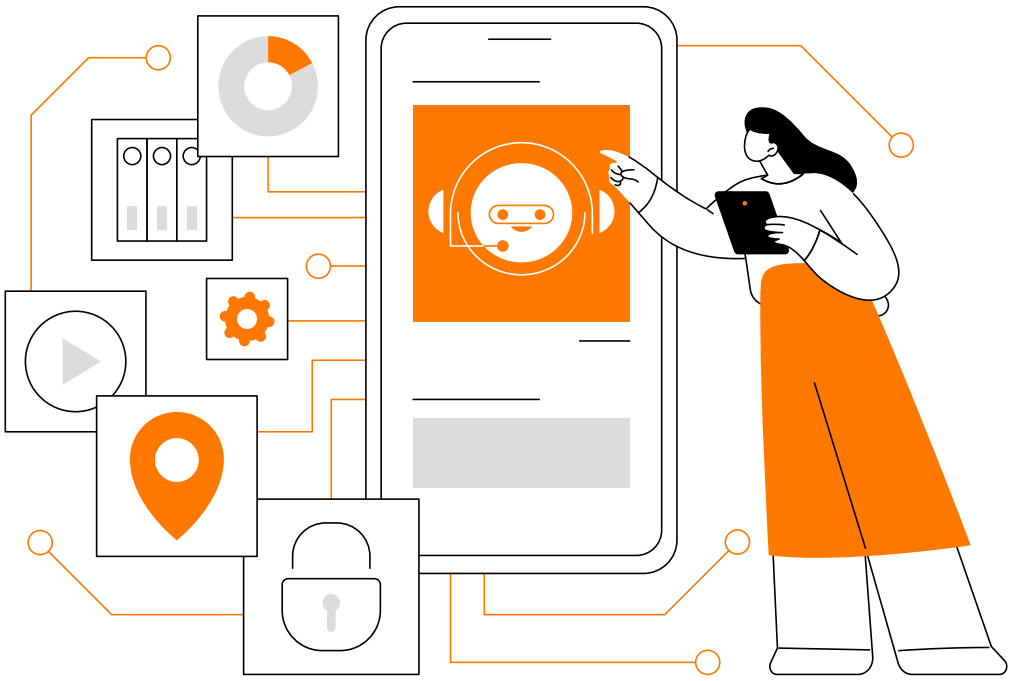

Improving customer relations

Another challenge for AI in business is to improve customer exchanges as part of an omnichannel strategy (phone, social networks, web, print, store, etc.) through the LLMs' power to reproduce human language. Generative AI can create a more fluid and effective exchange than traditional chatbots based on a specific pathway using all the data collected from customer exchanges. The bot will thus assist the sales force in managing customer relations, which is synonymous with improved customer satisfaction, for example, by reiterating a certain number of actions to be carried out, summarising the key points of previous customer contact points, or providing advice and recommendations.

Improving decision-making

Data analysis tools aim to support decision-making. For example, suppose we continue our illustration in Human Resources. In that case, it is possible to analyse data relating to headcount, payroll, carbon footprint, skills requirements, etc. and make policy, recruitment, training, or geographical mobility decisions to adapt to local needs. With Generative AI, it will be possible to enrich decisions by integrating third-party information from the web. This capability brings diversity and, above all, puts decisions into a global context.

A new source of innovation

With new business models, new ideas, new products, new services, and molecules hitherto unknown in biology or pharmacology, Generative AI can rely on economic situations, competitor data, technical documents, and market trends, enabling companies to diversify and stand out from the crowd. In this sense, Generative AI brings a broader perspective than just the company's perimeter and, when used properly, is a natural source of competitiveness.

The benefits of Generative AI

While the direct benefits of innovation are more complicated to calculate, the benefits of automation are more apparent. The productivity gains brought by Generative AI enable companies to optimise their costs and improve their presence with customers online.
Productivity gains and financial gains are converging: the time saved through automation will generate clear financial gains by replacing manual tasks and making it possible to carry out new tasks that were previously impossible. This is also what will enable us to innovate faster and further.
A Stanford University study reveals that using Generative AI could increase business productivity by 14%. According to MIT researchers, this increase could even reach 40%, while quality would be improved by 20%[12].
Another advantage is talent retention. In fact, by enabling teams to dedicate themselves to higher value-added and, therefore, more rewarding tasks, Generative AI contributes to well-being at work and strengthens employee commitment. It's a way of adding value that gives meaning to work and enables everyone to achieve more daily.

A Stanford University study reveals that using Generative AI could increase business productivity by

14%

Current limits on Generative AI development

Like all new topics, it is still challenging to master, which limits its use and benefits.

1st hurdle: AI hallucinations

The principle of Generative AI is to predict the next word based on the context and historical words it holds. So, as it is subject to bias and false data, AI is not 100% reliable and can come up with false answers. Consequently, efficiency gains cannot be guaranteed, and human review of the generated content is always necessary. In our HR example, we need to check that the AI has not added an unnecessary or unjustified clause to the employment contract or one that could even put the company at risk.

2nd hurdle: The risk of confidential data leakage

Some companies, such as Samsung, have banned their employees from using Generative AI tools such as ChatGPT or Google Bard. And with good reason: three employees of the semiconductor division entered sensitive company information into ChatGPT to assist them in their work searching for flaws in a database, optimising lines of code or summarising a work meeting. As a result, this information ended up in the database used to train the AI. This example illustrates one of the weak points of Generative AI: data security and ownership.

3rd hurdle: Loss of knowledge

The other indirect side-effect is the dependence on Generative AI. Some teachers fear that students will use a tool like ChatGPT intensively and are concerned about a possible loss of knowledge. Why should you bother to learn when you have a solution that has the answer to everything (or almost everything)? We must be cautious about the convenience offered by Generative AI because artificial intelligence is not intended to replace human beings. Taking a step back and challenging the information provided is still necessary. Generative AI offers genuine assistance through automation. It does not replace reflection. In the end, the decision remains human!

4th hurdle: Embracing artificial intelligence

An interesting comparison can be drawn between the calculator and Generative AI. In 1966, some mathematics teachers in the United States protested using calculators during lessons, fearing that students would no longer learn how to calculate. But we can still do mathematics and solve complex equations (except that we no longer need to know them by heart). So, we need to ask ourselves what limits we should set to provide practical help and best use human capacities.

How can we overcome the limitations of Generative AI?

As with any IT project, the introduction of Generative AI in companies must be supervised and supported. A framework must be defined for the use and maintenance of these AIs over time so that they can be monitored, made reliable and controlled. This control lets us know how the data is used to avoid bias and ensure that the solution always provides the right level of response or information to users. The aim is to ensure that the solutions deployed remain effective over time and to assess their contribution. Finally, the most important thing is to ensure that teams are acculturated and supported in their use so that they can make the most of Generative AI. This support has the advantage of reassuring users about the long-term viability of their business and encouraging awareness of the limits of these tools. As with any transformation, the key factor in the success and adoption of a tool is, above all, human.

Expert opinion

Samia Bouchareb, AI Ethics Officer, Orange Business

"At Orange Business, we are very attentive to the 'human control' aspect, which is also one of the seven key principles of the future European regulations. This stage of checking the Generative AI results is essential before using it in a department or generating a document. We always ensure that the results are verified by human input. The idea is to ensure that people always retain control over the AI system. That's why we use the models mainly internally. We are still in the experimentation phase, with a rather restrictive approach that we will open as we gain maturity and have more experience with the real risks encountered. We're likely to discover new risks in practice that we hadn't anticipated."

12 Generative AI at Work, Erik Brynjolfsson, Danielle Li & Lindsey R. Raymond, April 2023 Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence, Shakked Noy et Whitney Zhang mars 2023

Focus on business use cases

How can Generative AI be used in business? What new business opportunities does the technology open? On the other side, what are its limits? Like any new solution, Generative AI raises questions and concerns because it creates new ways of doing things, as you can see from the following examples related to content generation.
While many applications are still to be developed around Generative AI algorithms, it is already possible to exploit Generative AI for concrete use cases, providing added value for businesses.
Automatic content creation tools have been around for several years. In this sense, Generative AI has enabled them to evolve and, consequently, to improve.

“It is already possible to exploit Generative AI for concrete use cases, providing added value for businesses.”

Use case #1

The evolution of chatbots and voice bots in automated customer service scenarios

Chatbots are a clear example of the developments brought by Generative AI. Conversational robots are now an integral part of our daily lives, and they have not waited for Generative AI to be used in many applications. The best known of these are virtual assistants capable of directing online customers to the right product, page, information, or point of contact. Generative AI, on the other hand, will bring an extra level of fluidity and relevance to the response.
Until now, bots’ answers were formatted, and we often switched to human contact. Thanks to its training method, based on many sources of information, Generative AI will integrate much more complex, high-performance models, enabling more detailed processing of questions and, by extension, better quality answers. Companies will, therefore, be able to take the quality of their chatbots and voice bots to a new level. The result will be a better experience and greater customer satisfaction.

Use case #2

Creating a support centre or a knowledge database

Another application close to the chatbot example is the intranet, the dedicated agent, or the search engine. These engines often work by indexing documents or information but with more or less effective results. In this case, Generative AI could transform all searches into user-friendly natural language, consolidating information to provide a more precise answer to the query.
In practical terms, let's take the case of human resources. By integrating all HR information into the intranet, employees could ask questions and automatically access the correct document or the information they require, with a 'natural' response similar to an HR colleague.
The same applies to specialised agents who have been given a knowledge base to learn. In the case of a scientific thesis, for example, the integrated engine could use keywords to access the relevant section directly or provide a precise summary of a given topic or section.
The development of Generative AI means that the search engine can be used in the same way as Google. Instead of entering keywords and browsing several pages of answers, the Generative AI could summarise the few most interesting and relevant pages and propose an easy-to-read summary answer. We are, therefore, using a search engine by name, in other words, improved by an agent with an extensive knowledge base.

Use case #3

Automatic code generation

While there are many business use cases, Generative AI also provides invaluable assistance to developers by automating code generation, which saves a significant amount of time. The principle is the same: all you must do is enter a problem in text form, in natural language, and the tool can produce several hundred lines of extremely well-constructed code.

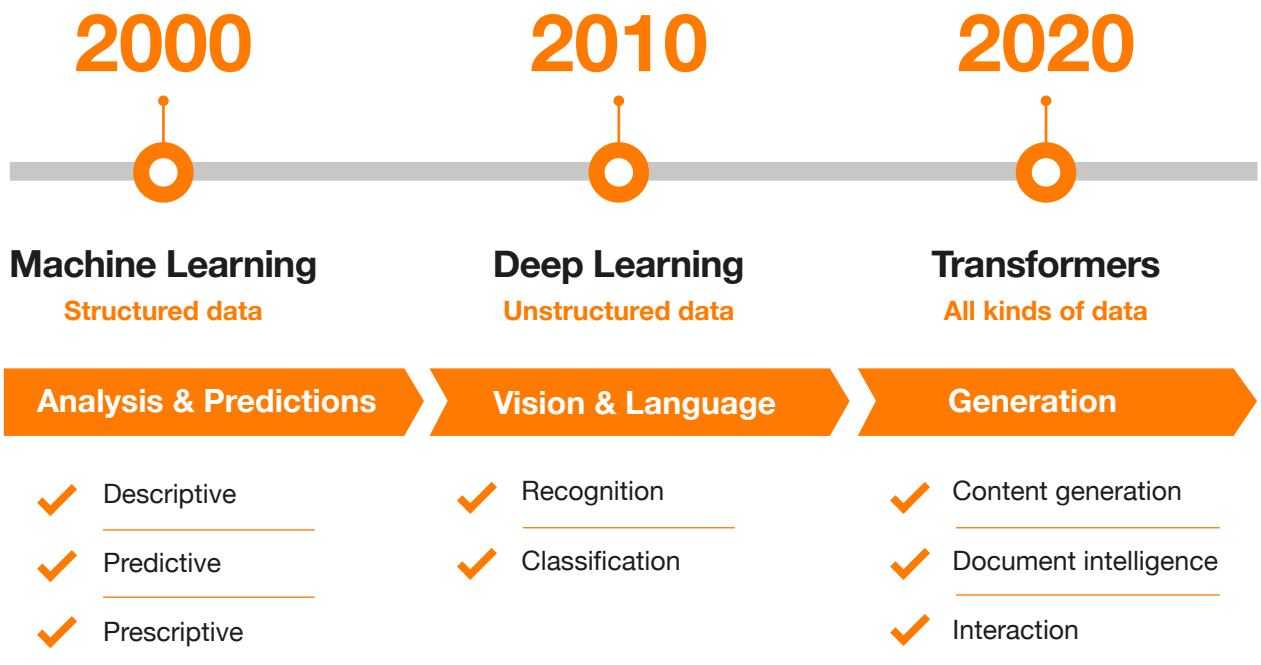

Unlike “classic” code generation, Generative AI brings an interactive dimension and an unequalled speed. Although it is now possible to find online tools that can perform all these tasks, the big difference with Generative AI is that it will transform “pretty much what I’m looking for” into “exactly what I want”.

Use case #4
Data generation

Generative AI can, therefore, be applied to many business use cases. But let’s be careful: it’s essential to understand that what makes Generative AI so compelling is the “extraordinary” amount of data used to train it. Chatbots, for example, need information to provide correct and relevant responses.
So, data is an absolute prerequisite for developing business use cases for Generative AI. Without data, there can be no data project and no AI!
Nevertheless, companies often have more or less large data samples without falling into an extreme case where no data is available. Not only would Generative AI be able to learn from this data, but it would also be able to generate a certain amount of additional information. The idea here is to summarise additional and varied data from a sample to improve machine learning models. It then becomes possible to create models with very few parameters.
For example, when developing software, it is necessary to use test data before going into production. The data created by Generative AI would be ideally suited to these tests.

Use case #5
Creative content creation

Lastly, this final use case is based on the creative capabilities of Generative AI. Although the primary use cases are based on text-based AI for information purposes, AI is also capable of a specific form of creativity. The marketing and communication fields are particularly appropriate for this, for

	IT	Marketing/ sales	HR/legal	Others
Documentary intelligence	▪ Development assistance ▪ Documentary proofreading	▪ Automation of routine sales tasks ▪ Analysis & summary of customer feedback	▪ HR Assistant ▪ Legal compliance control, contract analysis	▪ Summary of meeting ▪ Search for product feedback, internal topic
Content generation	▪ Code & documentation generation ▪ Test data generation	▪ Content creation ▪ Creating content for social media	▪ Newsletter content ▪ Job offer generation	▪ Generation of web interface ▪ Creation of products, molecules, materials
Interactions	▪ Generation of test scenarios ▪ Create simulation environments	▪ Campaign Marketing generation ▪ Chatbot, FAQ, products	▪ Virtual trainer, Onboarding assistant ▪ Drafting simple contractual clauses	▪ Suggested answers to customer support tickets ▪ Generation of virtual property tours

“In some cases, Generative AI will bring a simple evolution. In others, it will enable new uses to be created or more advanced performance to be achieved.”

example, through the generation of images to illustrate an article. In this case, Generative AI provides an additional visual source that meets the rapid response needs. Some projects require urgent creation. Marketing teams could, therefore, ask their content generation tool to create a graphic that they could, where necessary, rework later, gaining precious time in the process. Tools such as DALL-E (also offered by OpenAI), Midjourney and Stable Diffusion can already be used to create relevant images from detailed text descriptions.
In a completely different field, Generative AI enables further simulation of human behaviour. Like video games, developing a professional use case in training or remote control of certain industrial activities is possible. Why not consider an augmented or virtual reality helmet capable of precisely reproducing the human gestures to be performed in different scenarios?
These tools are new, and many uses are still in the pipeline. As a result, we have minimal experience of the results obtained. This lack of feedback means that many companies are reluctant to embark on the adventure of content generation. How far can we go? What limits will the production of information generate? These compelling technologies are highly present in the cloud. So, the question of data ingestion and confidentiality becomes more pertinent. But we are in a state of permanent evolution because AI is already changing the way we work. In some cases, Generative AI will bring a simple evolution. In others, it will enable new uses to be created or more advanced performance to be achieved. This is where Data and AI experts such as Orange Business will have a fundamental role to play in promoting understanding and awareness since large quantities and high-quality data are needed to obtain a response quality comparable to that generated by mass-market solutions such as ChatGPT.

Expert opinion

Éric Blanche, Digital Offer Lead, Digital Services France, Orange Business

“In the current era, customer relationship centres are facing an exponential growth in interactions, seeking the perfect balance between industrialisation and personalisation, without neglecting the economic and customer impact. Verint, the leader in customer engagement platforms with over 10,000 customers, is transforming the call centre environment by integrating Generative AI into its solutions, making it easier to summarise interviews, optimise the knowledge base and analyse data.
Far from being limited to strictly following a script, Generative AI now makes it possible to analyse the stylistics and intent of the operator, bringing a human touch to interactions. This innovation gives operators greater freedom, improving the fluidity and efficiency of the conversation while saving time for the agent, both during and after the interaction.”

Adopting Generative AI: The Right Methodology

74% of executives believe that the benefits of Generative AI will vastly surpass the concerns raised.[13] However, only 19% of companies have launched concrete initiatives on the subject.[14] Where to start? We have summarised the key steps in this section. What are the requirements for Generative AI? What are the key questions to ask beforehand? Even for professional Data Scientists, it is sometimes challenging to embark on an artificial intelligence project, given that the technology is causing companies to rethink how they work. Generative AI is no exception. There are three essential pillars for launching the development of a use case.

Step #1
Don't start from scratch!

Today, several pre-trained language models (LLMs) already make it possible to process a certain number of objectives using a conversational assistant right from the start of a project. The same applies to image recognition. It is possible, and even recommended, to use resources made available by the community, whether through open-source algorithms or publishers such as Microsoft with Azure. Teams already have models capable of recognising an

Acculturation

AI Insiders

- Webinars
- Seminar
- Workshop on AI Acculturation

Ideation

AI Ideate

Workshop to identify use cases

image, object, verb, or adjective in a sentence... In short, they save a great deal of time! In the event of a shortage of available resources, it is in the teams' interest to seek support in advance to identify and test existing solutions they can use without reinventing the wheel!

Step #2
Expand the boundary progressively

A Generative AI project is not so different from a traditional artificial intelligence or data science project. As a result, if too many functionalities or complexity are included from the start, the objective will be too ambitious and, therefore, difficult to achieve. In addition to taking an excessive amount of time to develop, the main risk is that you will have to deal with many simultaneous problems, leading to the project's failure. For example, the solution is to gradually develop a conversational assistant: start with an assistant that can only handle 20% of the problems or use cases, then slowly improve it! Agility and the ability to expand the scope as you go along are essential success factors in a Generative AI approach!

Agility and the ability to expand the scope as you go along are essential success factors in a Generative AI approach!

It is essential, at the start of the project, to clarify with the business the range of functions that you wish to integrate into the conversational assistant and to classify them by block of functionalities or by block of use cases defined along two axes:

- A business axis established by future users
- And a technical axis on which the experts are trying to quantify the algorithmic complexity

The objective is to start by integrating "quick wins", i.e. those with high business value and low technical complexity. Then, based on user feedback, the assistant is enriched at each sprint with new functionalities, new use cases, new databases, etc.
In this iterative approach, we estimate each time the load of each source in terms of technical integration and business interest.

Assessment
AI Roadmap

- Complete assessment of Data & AI Maturity,
- Deep dive into AI ambitions
- Definition & prioritisation of use cases

Poc
AI LAB

POC creation, scalable by design, of AI solution in a few weeks

Scale
AI Scale

Industrialisation & deployment at scale

Change
AI CoE

- Guidance
- Targeted coaching sessions

Step #3

Test and display your algorithm

Not only do you have to test your algorithm against the business, but you also have to do it as quickly as possible. Please don't wait until a product is perfect before presenting it to the business! This sharing requires the ability to evolve soon, implying a certain flexibility to automate the delivery and learning of the different models and, ultimately, to enhance the functionalities.
The scalability of the infrastructure may also be an essential factor in the methodology because all this data can represent large, even unprecedented, learning volumes. It depends on whether Generative AI integrates only internal, "open data", or other external data. This is why having a scalable infrastructure, or at least one that can be easily upgraded, is also a critical success factor.

How can we facilitate the adoption of Generative AI?

We started with some of our customers' use cases to answer this question. Here's how Generative AI has helped them solve problems and fast-track their projects. Business workshops will be necessary to identify and retain the proper use cases to develop. These workshops will be even more productive if acculturation for the different business stakeholders precedes them.

Case #1

You want to develop your algorithm

Using Generative AI allows you to accelerate your labelling phase (image, text, etc.). This is a significant advantage, as data labelling represents a non-negligible part of computer vision and NLP projects. Today's tools make it possible to recover significant quantities of labelled data. The key is the ability to learn your models based on this data by controlling your algorithm end-to-end.

Case #2

You don't have enough data to teach your LLM

Use Generative AI to build up a consistent learning base. In fact, why not ask GPT (or others) to generate a thousand different ways of formulating a question, for example, using different languages and styles? The accelerating nature of Generative AI means you can make your model learn much more quickly and efficiently.

"Business workshops will be necessary to identify and retain the proper use cases to develop."

Case #3

Can a partially private / partially open tool be created?

The question arises in cases where particular functions are required. In this case, developing an in-house conversational assistant will sometimes be necessary. On the other hand, for more generic issues, it is perfectly possible to use a more general solution, such as Azure OpenAI. The best approach is always to start by educating the business on what Generative AI can (and cannot) do in the company's context, followed by ideation workshops before launching into development. As with any AI project, the methodology to be adopted is fundamental to the project's success.

How can we implement a progressive Generative AI approach?

Just as a reminder, Generative AI enriches the model little by little with increasingly relevant functionalities and use cases. A process which, contrary to what you might think, is not so simple.

And why is that? It is not recommended to use a tool like ChatGPT in a company due to the problems of confidentiality and data security.

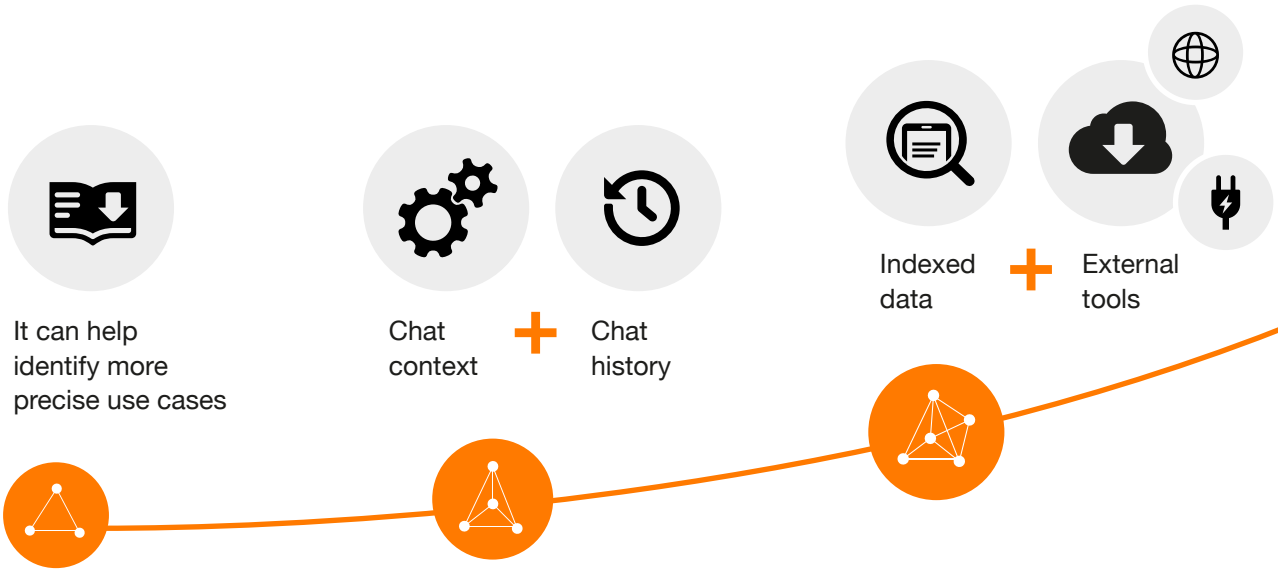

1 ‘Zero-shot’ mode

Deployment of a generic LLM model

The project’s first stage involves deploying an equivalent of ChatGPT at the enterprise level, using a tool that respects confidentiality. This will generally be a Large Language Model (LLM) already pre-configured for the company (for example, Azure Open AI). All you must do is take this LLM and make it available to your employees. This is known as “zero-shot” mode. It involves taking a generic LLM model, whatever it may be, and putting it in secure environments for the company or even in the cloud, but safely.

2 ‘Few-shots’ mode

Conversational bot that indicates its context & with its history

The conversational bot is based on a context and has a prompt record. It is, therefore, used for professional purposes. “Few-shots” mode involves refining the assistant’s responses. Each time the model is queried, it is given as many contextual elements as possible, so-called prompt engineering, concerning the expected response.
The answers given will, therefore, be better adapted to the company’s vocabulary and problems.

3 The agents

In addition to the context add memory (indexing) & actions (external tools)

In this next step, these models must be interfaced with the company’s operations to be used easily without disrupting processes. In addition to the context, source documents (indexing) and actions (external tools) are added. The assistant will then go through the indexed input data and the connection with the external tools to expose the responses via an API, for example.

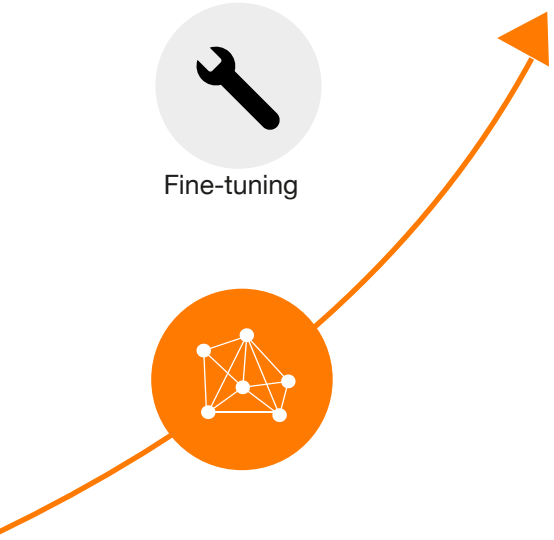

4 Training / Re-training

Training & enrichment the model to specialise in (Fine-tuning or Lora).

The final step is to train, re-train, or over-train these general AIs with the company's entire language. This generally involves customising the last supervised layers of the model. This is how Generative AI will gradually create value within the company.

Generative AI remains, first and foremost, an artificial intelligence project, mainly through the project development phases. Like many AI projects, it can be highly complex. To get started, it is therefore advisable to call on experts in the field, to get support, and to take the time to consider the pros and cons of each solution.

Expert opinion

Michael Deheneffe, Strategy, Marketing & Innovation Director, Digital Services France, Orange Business

"One method of limiting the negative effects or potential pitfalls of Generative AI could be to set up a feedback loop. In other words, in the case of content generated by an AI, the role of a supervisor will be to check whether the production is correct or not, whether the text is of good quality or not, and so on. A 2nd option is based on the ability of the Generative AI to indicate the sources it has used, which will make it possible, as in any documentary work, to cross-check the sources and verify the reliability of each of them. This traceability of Generative AI should make it possible to correct some of the errors linked to data quality. Finally, Generative AI requires training! Any tool, whatever its nature, has both positive points and limitations. That's why it's important to have the necessary culture to measure the impact, the risks, the limits, the inappropriate behaviours etc., of Generative AI. So, whether taken across society, at the company level or the personal level, it is essential to train a maximum number of people to understand what Generative AI is, what its dangers and limits are, and to deploy a sufficiently high level of vision to detect any biases. Generative AI will only work if humans are even more human."

13 According to the Capgemini Research Institute report: Harnessing the value of Generative AI: Top use cases across industries, 2023
14 Gartner Poll Finds 45% of Executives Say ChatGPT Has Prompted an Increase in AI Investment

Deploying trustworthy AI

68% of employees use ChatGPT at work behind their employer's back.[15] This means more than two out of three companies share personal or confidential data with an unprotected public application. This alarming finding provides an opportunity to raise the question of the ethical and legal considerations of Generative AI and, more broadly, of trusted AI. What risks could the uncontrolled use of Generative AI entail for businesses? Let's look at some possible answers.

The question of AI ethics goes far beyond the company's walls. It covers a wide range of societal issues, such as the risks of disinformation, standardisation of thought, the social damage caused by click workers and ecological issues. While these issues are at the heart of the debate around Generative AI, it is impossible to deal with them all here. This chapter, therefore, proposes to focus on the ethical issues to be considered by companies when deploying Generative AI applications.

What is AI ethics?

Ethics are intrinsically linked to the values and culture of each human group. Determining whether an AI is "good" or "bad" is impossible. Indeed, this perception is inevitably linked to the culture of the country in which it is applied, as well as to the experience of each individual and many other personal factors.

This is at the heart of this subject's complexity, which is becoming highly philosophical!
For AI applications in business, the ethical question will have to be approached from the perspective of acceptability. Is the application of AI for a given use case acceptable for a group of users?
What are the conditions of deployment to make it acceptable? What limits and safeguards need to be deployed to make it acceptable? All these questions need to be asked by applying them to the group of users or people exposed to the uses of AI.

Why ask the question of trusted AI?

Confidence in AI, and by extension in Generative AI, is being blurred by the crucial issues of ethics and legal liability, combined with irrational fears that are being drained by science fiction and the media, which are making headlines with sensationalist announcements that are not very serious.
While it is essential to identify the critical issues, the number of incidents linked to unethical abuses of AI is exploding. According to the AIAAIC database, which lists them worldwide, the number of AI-related incidents and controversies has multiplied by 26 between 2012 and 2021,[16] rising from around ten to 260 recorded incidents. These include numerous deepfakes and problems linked to the biases included in the training data.
With the risks to exposed individuals becoming increasingly numerous and tangible, companies must take this issue seriously and make it a core part of their Generative AI projects.

The AI Act, the European foundation for trusted AI

Like the General Regulation on the Protection of Personal Data (RGPD), Europe aspires to be a pioneer in regulating the uses of AI. It has, therefore, been working for several years on an ambitious draft regulation validated by the European Parliament on 14 June 2023, paving the way for implementing a legal framework for AI uses.
Companies will thus have to classify all their uses of AI according to their level of ethical risk: minimal, limited, high and unacceptable.
With this regulation, Europe sets a precise line about uses with a very high social and societal impact. It is setting itself apart from other regions, such as China, which extensively uses facial recognition or social rating.

"Compliance with the AI Act can start now by identifying and classifying all uses throughout the company."

The company must combine appropriate risk reduction measures depending on the level assessed for each AI application. Therefore, compliance with the AI Act can start now by identifying and classifying all uses throughout the company. The principles of ethics-by-design (or ethics-by-construction) should be introduced to consider these issues at every project stage. Our white papers place them at the heart of the AI construction and deployment methodology.

Generative AI & limited risk

Transparency, copyright respect

Chatbots, avatars, Generative AI applications in general.

High risk

Compliance review

Education, employment, justice, immigration, law, biometric identification, critical infrastructure, management public & private services

Minimal risk

Code of conduct

Anti-spam filters, video games, scenarios

Unacceptable risk

Prohibited

Social scoring, facial recognition, manipulation, dark-pattern interface (dark-pattern AI)

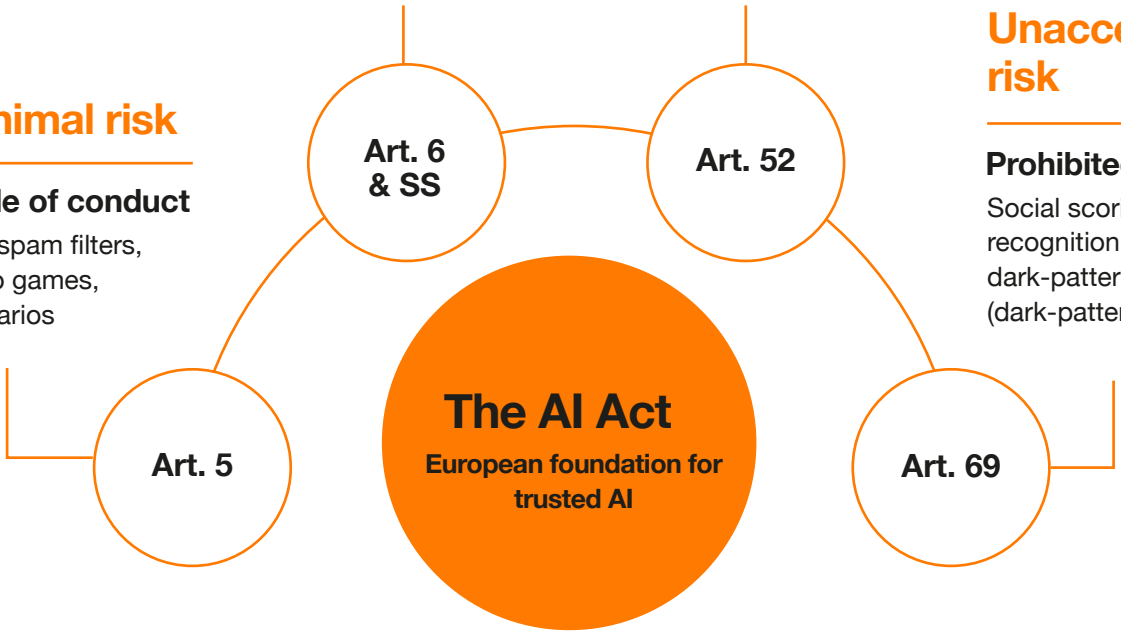

The 3 pitfalls of AI

The way AI operates, based on machine learning, presents inherent problems. There are no apparent solutions to these problems, but they must be understood and assessed to promote trusted algorithms.

1 The issue of fairness

When we talk about training data, we inevitably talk about the risk of bias, particularly given the phenomenal volume of data needed to train Generative AI models. The list of sites on which ChatGPT has been trained[17] clearly illustrates the lack of cultural diversity and representativeness of populations or schools of thought. For example, the American start-up Textio highlighted multiple social biases[18] (linked to age, gender, or ethnic origin) during a test of 167 job ads written by ChatGPT.
Algorithmic biases are thus closely linked to the lack of diversity and inclusion in the databases used. To develop AI that can be trusted, it is necessary (but not necessarily a sufficient condition) to carry out in-depth statistical tests on the training data and to increase diversity and inclusion in tech to broaden the profiles of AI developers: gender mix, cultural diversity, different educational backgrounds, social openness, and so on.

2 The transparency challenge

When it comes to transparency, one fundamental principle must be respected. We must never create confusion about who speaks in a human and an AI dialogue. Users must always know that they are dealing with an AI. This duty to inform is essential and forms the basis of the principle of transparency of use.
Another problem is the lack of clarity in the decisions and content generated by AIs. Since these AIs are based on transform-type models with extensive and complex neural networks, this makes them opaque and complicates the detection of their “hallucinations” (false answers presented so plausibly that they can be considered valid).
This is why companies need to be able to explain why and how AI makes its decisions. And how? By specifying, for example, what sources of information have been used to compose the answer. This notion of the explainability of AI decisions is essential to comply with ethical principles and is integral to AI transparency issues.

3 The accountability challenge

The company remains fully responsible for the AI it deploys, from a legal point of view, and consequently for the decisions of its AI solutions. It, therefore, accepts the risk of amplifying an ethical problem linked to AI’s capacity for extreme automation. This is why the sooner it integrates the moral dimension into the development of its AI, the more it will limit the risks.
As we can see, the subject of ethics in AI covers several areas: Data protection and, at the very least, compliance with the GDPR is essential. Without protecting personal data, we cannot claim to establish trusted AI. Then, we need to look at the structural properties of artificial intelligence, particularly those of fairness and transparency, which need to be considered by the data scientists who set up the algorithms. Finally, ethics is ultimately a philosophical issue. Each company will have to choose the fields of application and uses and define the limits and safeguards to be implemented.

Points to consider when establishing a trusted AI:

Compliance with regulations & ethics

- Philosophy & Culture
- Choice of uses
- Limits & safeguards

Structural properties of AI

- Fairness
- Transparency
- Accountability

Data protection

- Consent
- Security
- Governance

The environmental impact of AI

Another issue to consider when developing respectful AI is its environmental footprint.
The training and operation of Generative AI involve significant resources for data storage and algorithmic calculations, which inevitably have a high environmental impact.
Hence, it is essential to use pre-trained algorithms systematically (GPT stands for Generative Pre-trained transformers) and adapt them to the company's specific use cases (using transfer learning, for example). This avoids re-training billions of pieces of data for each enterprise solution. The enormous environmental cost of training generic AI (such as GPT, Bard and others) is then diluted to serve millions of uses on a global scale. In addition, several algorithmic approaches are currently being developed that could lead to equally effective algorithms with less impact on the environment.

Frugal AI and eco-design of AI algorithms

The concept of frugality was launched in 2019 through the advocacy of researchers at the Allen Institute for AI in favour of frugal AI or green AI, i.e., AI that is more efficient, inclusive, and less greedy regarding data and energy (linked to computing power).
However, the astronomical amount of data used to train Generative AI, such as ChatGPT, is helping to increase the carbon footprint of digital technology. Companies need to integrate the concept of frugality into their processes to deploy virtuous AI that consumes less energy.
This is why, at Orange Business, eco-design principles are applied to the methodologies deployed.

Expert opinion

Samia Bouchareb, AI Ethics Officer, Orange Business

"In response to the risks associated with Generative AI, Orange has drawn up an ethical charter for data and artificial intelligence.[20] The aim is to set out the company's moral values concerning data and artificial intelligence. The Orange Data and AI Ethics Council28 drew up this charter,[21] comprised 11 external figures chosen for their independence, neutrality, expertise, and the diversity of their profiles. Eventually, companies will also have to adopt a practical user guide. The aim is not to prohibit but to encourage responsible use. Technologies are vectors of operational efficiency, facilitation, and customer value. Hence, the idea of a charter, but above all, a practical user guide.

At Orange, we also provide our employees with a 45-minute training course accessible online to all, enabling them to become acculturated to the subject and fully aware of the risks associated with using AI. This acculturation phase is critical and needs to be as broad as possible.

The benefit today is that companies can build on their experience of the GDPR to facilitate the implementation of responsible AI 'by design' and the deployment of AI Ethics Officers in each entity to be as close as possible to the real-life uses of AI. We have set up a network of ethics officers within the Orange Group. Each central business unit has its own ethics officer, who is responsible for anticipating regulation, educating teams, setting up governance and overseeing the implementation of processes. But don't wait for the IA Act to be applied to launch this stage. At the same time, we have set up a multidisciplinary responsible AI committee within Orange Business, consisting of data scientists, data privacy experts, a legal representative, a data governance specialist, and so on. This complementarity of expertise is essential. The transformation generated by Generative AI is a collective and multidisciplinary issue."

What are the concrete actions?

Therefore, any Generative AI approach must consider multiple issues relating to trust, security and ethics. Whether to comply with AI Act regulations or to ensure the best possible use, all these issues must be placed at the heart of any AI deployment approach. 12 concrete actions should be prioritised for deploying trusted AI in companies.

6 actions that Data Scientists & Data & AI professionals can implement:

- Establish the principles of ethics by design;
- Profiling learning data to confirm its representativeness;
- Test and validate algorithms against ethical criteria;
- Audit AI in production according to ethical criteria;
- Adopt algorithmic sobriety and the principles of frugal AI;
- Sign the Data Scientist Hippocratic Oath (or equivalent).[19]

7 principles for ethical AI

According to the European Commission, seven elements are essential to achieve trustworthy AI:[22]

1. Human factor and human control
2. Robustness & safety
3. Privacy & data governance
4. Transparency
5. Diversity, non-discrimination & equity
6. Societal & environmental well-being
7. Accountability

15 70% of Workers Using ChatGPT Haven't Told Their Bosses: Survey (businessinsider.com)
16 AI Index report, Stanford University, 2023
17 Inside the secret list of websites that make AI like ChatGPT sound smart, The Washington Post, April 2023
18 Textio blog https://textio.com/blog/chatgpt-writes-job-posts/99089591200
19 https://towardsdatascience.com/hippocratic-oath-for-data-scientists-407d2db15a78?gi=dcea12a0d0f1
20 Charte éthique ENG.pdf (orange.com)
21 Data and AI Ethics Council: representing responsible AI | Corporate (orange.com)
22 To find out more, download our white paper "Ethical data / ethical AI: the 2 faces of a responsible future", December 2019

Conclusion

Towards a more human digital world

Generative AI opens a whole range of opportunities for businesses and will create new needs. In the future, we could see the development of data marketplaces to compensate for the lack of data and feed models. The media could also begin to monetise their knowledge bases and provide information on the "state of the world", updated in real-time. We will see trusted data referenced, valued, and monetised as data is strengthened.

Generative AI, the new digital must-have

More concretely, for businesses, Generative AI will eventually be integrated into all everyday digital tools, opening a whole new world of possibilities. For example, AI included in communication and collaboration tools can take minutes of meetings, summarise critical points, define a list of tasks to be carried out, schedule them in the diaries of the people involved and even send reminders if people forget. In this sense, Generative AI will "complete" and assist the employee by taking on the least rewarding and most time-consuming tasks.
"AI will enable us to move towards what we call the augmented employee", says Samia Bouchareb, AI Ethics Officer at Orange Business. HR will have a vital role in this transformation of businesses to anticipate and support these changes in professions. Generative AI will also raise the question of how to integrate it into products to generate greater customer satisfaction or a better user experience. While it is still too early to know what type of AI to implement and/or market, the subject is already strategic for companies. Companies that don't position themselves quickly on this issue will fall far behind in competitiveness."

Accompanying rather than opposing

There's no point in opposing Generative AI because it's already here, within companies themselves, sometimes as a partner and sometimes as a competitor. That's why, far from prohibiting it, we need to support it. The question is not whether it should be deployed internally but how it should be deployed and, above all, for what purpose. This last point is key. Businesses, particularly in Europe, will have to ensure that Generative AI is used for laudable purposes compatible with their values and future regulations. This represents a major strategic challenge for the organisation's external ecosystem and its internal stakeholders.
"Tomorrow, the use of technological tools such as Generative AI will be considered by customers and employees when choosing the partners they wish to work with, just as social and environmental responsibility may be. As a result, companies need to start thinking about and testing Generative AI now. They need to ask themselves how they can master artificial intelligence, specifically Generative AI, before artificial intelligence masters their own business by disrupting their market when they are not ready", adds Michael Deheneffe, Strategy, Marketing & Innovation Director, Digital Services France at Orange Business.

In transformation, there is "formation"

That's why thinking ahead today is essential to avoid being left behind tomorrow. This necessity implies massive training in using this Generative AI on universities and companies' benches. How can this be done? By upskilling or retraining the employees most interested in AI and data - statisticians, engineers, computer scientists, data engineers and even data scientists. This is why, against a backdrop of a shortage of IT profiles, there is an urgent need to train people who have already mastered the company's processes, tools, businesses, and legacy so that we don't have to wait for the new generation to graduate before we can make the most of Generative AI. Acculturation is synonymous with attracting and retaining talent!
For Michael Deheneffe, expert companies such as Orange Business have "from a societal point of view, the major challenge of helping all generations to master these technologies which will arrive very quickly".
"On the other hand, one essential point needs to be included in this training: taking account of the ethical, economic, and social aspects of Generative AI, in other words, the philosophy applied to businesses, or at least applied to the economy and society. In addition to courses in coding, data management or computer engineering, classes in sociology, demography, philosophy, ethics, etc. will be crucial", concludes Michael Deheneffe. Tomorrow, we will have ethicists specialising in AI who will measure, manage, guide and train algorithms to perform tasks that correspond to the company's values. But tomorrow is being shaped as of today! The Generative AI revolution is arriving at high speed and beware of companies that are too slow to jump on board!

Contributors

This white paper is directed by Didier Gaultier, DataScience & Ethics AI Director, Digital Services France, Orange Business.

Authors of chapters, articles, and boxes (in order):

- **Aliette Mousnier-Lompré:** General Manager Orange Business
- **Didier Gaultier:** Director DataScience & Ethics AI, Digital Services France, Orange Business
- **Paul Ferré:** Lead Data Scientist & Senior DataScience Consultant, Digital Services France, Orange Business
- **Romain Bernard:** Lead Data Scientist & Senior Manager Data Science, Digital Services France, Orange Business
- **Nathanael Guiberteau:** Senior Manager, Digital Services France, Orange Business
- **Pierre-Henri Mas:** Lead Data Scientist & Senior DataScience Consultant, Digital Services France, Orange Business
- **Erwan Josse:** Lead Data Scientist & Consultant DataScience Senior, Digital Services France, Orange Business
- **Mick Levy:** Innovation Business Director, Digital Services France, Orange Business
- **Michael Deheneffe:** Director Strategy, Marketing and Innovation, Digital Services France, Orange Business
- **Éric Blanche:** Digital Offer Lead, Digital Services France, Orange Business
- **Samia Bouchareb:** Head of AI Ethics, Orange Business

Our special thanks for their participation, encouragement, and valuable comments:

- **François Fleutiaux:** Director, Orange Business France
- **Laurent Sicart:** Director, Digital Services France, Orange Business
- **Sandrine Le Seac'h Godard:** Marketing Director, Digital Services France, Orange Business
- **Safa Baghai:** Communication Manager, Digital Services France, Orange Business
- **Dominique Roda:** Communications Manager, Orange Business
- **Nassima Auvray:** Chief Trust Officer, Orange Business
- **Vincent Garnier:** Head of Data & AI, Orange Business
- **Audrey Lompré:** Marketing Manager, Digital Services France, Orange Business
- **Aurélie Goncalves:** Manager Innovation & Strategy, Digital Services France, Orange Business
- **Françoise Soulié-Fogelman:** Scientific Advisor, Hub FranceIA
- **Yann Le Cun:** VP and Chief AI Scientist, Meta

Acknowledgments:

- **Renaldo Candreva**: Graphic Designer

Data Science and AI Glossary

Agent: in artificial intelligence, an intelligent agent (AI) is a hardware and/or software entity with a certain degree of autonomy, capable of interacting with the environment to achieve objectives

AI Act: regulatory framework proposed by the European Union to regulate artificial intelligence applications.

AIAAIC: independent public interest association advocating transparency and openness in the field of AI, algorithms, and automation

AI winter: a prolonged period of reduced funding and interest in AI research, generally due to unmet expectations.

Algorithm: set of clearly defined rules and steps for solving a problem or performing a task.

Application Program Interface - API: a set of rules and protocols that enable different software programs to communicate with each other.

Artificial Intelligence - AI: a field of computer science focused on creating agents capable of analysing data, learning from it and interacting with the environment to perform specific tasks in a given environment, with a certain degree of autonomy. It is a branch of Data Science and aims to enhance human capabilities by facilitating decision-making and the execution of complex functions.

Augmented reality: technology that superimposes digital information (such as images and videos) in the real world, usually through a device like a smartphone or augmented reality glasses.

Attention mechanism: this is an advanced technique in natural language processing (NLP) that enables the model to assess the relative importance of each word in a sentence, depending on its position and context, when performing tasks such as machine translation or text generation.

AWS: Amazon Web Services, a cloud platform offered by Amazon, featuring compute, storage and machine learning services.

Bandwidth: a measure of the data transmission capacity of a network or Internet connection, generally expressed in bits per second.

BERT: a bidirectional language model based on transformers developed by Google, pre-trained on a large quantity of text.

Conversational agent - Chatbot: computer programme capable of conversing with a user in natural language, generally used to provide information or carry out tasks.

ChatGPT: chatbot developed by OpenAI, based on GPT architecture, capable of analysing and generating text in natural language.

Data Engineer: expert in designing, building, and maintaining systems and infrastructures to collect, store and analyse data, including creating and industrialisation data pipelines to feed and deploy AI models.

Data Science: interdisciplinary field that employs scientific methods and processes to build models and extract knowledge from structured and unstructured data, encompassing disciplines such as statistics, machine learning and artificial intelligence.

Data Scientist: an expert whose role is to understand business needs and find appropriate solutions by designing and implementing models based on data engineering, statistics, mathematics, machine learning and artificial intelligence.

Dall-E: AI model developed by OpenAI, capable of generating images from text descriptions.

Deep learning: a particular form of machine learning that uses multilayer neural networks inspired by the animal visual cortex to analyse various levels of data representation, particularly unstructured data such as images.

Deep neural networks: see Deep Learning, neural networks with many layers, enabling analysis of increasingly complex levels of data representation.

Exa (byte): unit of measurement for data quantity, equivalent to one trillion (1018) bytes or one thousand terms.

Facial recognition: technology that uses AI algorithms to identify or recognise a person based on their face captured in an image or video.

Few shots mode: machine learning technique where a model can learn and perform a task with limited training examples and little context.

GAN: a class of AI models used to generate new data that resembles a given training dataset, usually composed of two competing neural networks.

GDPR: General Data Protection Regulation, an EU regulation governing personal data protection.

Generative Artificial Intelligence / Generative AI: a branch of AI specialising in creating new data (such as text, images, videos, etc.) from instructions supplied in natural language, commonly known as “prompts”. These advanced models can generate content that closely matches the criteria specified in the prompt, facilitating fluid, natural interactions.

Google Cloud Platform (GCP): a suite of cloud computing services offered by Google, including computing, storage, and machine learning services.

GPT: Generative Pre-trained Transformer, a series of advanced language processing models developed by OpenAI.

GPT-4: currently the most advanced version of the GPT model (at the time of publication), with 175 billion parameters.

GPU: specialised processor, often integrated into a graphics card, designed to accelerate image and video processing, as well as to manage complex matrix or neural calculations, particularly in the context of training artificial neural networks.

Graphics memory: storage space in a graphics processing unit (GPU) to store temporary data and graphic images.

Hyperscaler: a company providing large-scale infrastructure for cloud computing services such as Amazon, Google, and Microsoft.

ImageNet: annotated image database widely used for training AI models in image recognition.

InstructGPT: a supervised layer of OpenAI's ChatGPT model, trained to follow text instructions and answer questions.

Large Language models - LLMs: AI models trained on large quantities of text and capable of generating text in natural language.

LLaMA2: language model developed by Meta, a technology company known for its social media and virtual reality developments.

Machine Learning: a branch of AI that enables machines to learn and improve from experience without being explicitly programmed.

Microsoft Azure: cloud computing services platform offered by Microsoft, featuring computing, storage, and machine learning services.

Midjourney: AI that generates images from text descriptions.

Neural Networks: AI models initially inspired by the functioning of the human brain, composed of interconnected nodes that perform statistical calculations on input data.

Nvidia H100 GPU: a specific GPU model developed by Nvidia, known for its high efficiency in processing AI tasks.

Open AI: OpenAI is an AI research and deployment company whose mission is to ensure that general artificial intelligence benefits humanity.

Open-source library: a set of functions that extend the capabilities of an existing programming language, such as Python, and are freely redistributable and modifiable.

Overfitting: a phenomenon where a complex machine learning model learns using too little training data, to the point of not generalising well on new data.

PaLM2: advanced language model developed by Google, designed to process, and analyse natural language.

Parameter: variable in an AI model that is adjusted during training to optimise model performance.

Reinforcement learning: a machine-learning method in which an agent strives to maximise a reward function by interacting with its environment. The agent can also adjust and improve through RLHF (Reinforcement Learning by Human Feedback), where humans provide feedback to achieve specific goals more precisely.

Self-supervised learning: is a machine learning method in which the model is trained to complete a missing or hidden part in an input (such as a word in a sentence) from a corpus of other inputs.

Sprint: project management method that divides work into small units called "sprints", generally used in Agile methodologies.

Stable Diffusion: AI that generates images from text descriptions.

Statistics: a branch of data science and mathematics that focuses on data collection, analysis, and interpretation. It aims to extract meaningful information from data to inform decision-making, forecast trends and explain various phenomena.

Supervised learning: a machine learning method in which a model learns to predict specific outputs from the inputs it is given using an output framework and a set of previously annotated examples.

Transfer Learning: a machine learning technique in which a model, initially trained on a specific task, is reused as a starting point for a secondary task. This method allows capitalising on the training already carried out with a learning complement to perform related tasks with less data and computing time.

Training (of a model): the process by which an AI model learns from training data to make predictions or decisions.

Transformer: artificial intelligence model architecture, mainly used in natural language processing. This structure, introduced by the publication "Attention is All You Need" by Vaswani et al., is based on the attention mechanism, which enables the model to assess the relative importance of each word in a sentence, depending on its position and context.

Unsupervised classification - typology: the process of categorising objects or information into groups or classes based on their common characteristics.

Unsupervised learning: a type of machine learning in which the model discovers links and patterns in the data without first being provided with a framework of answers and annotations, as in supervised learning.

Vector space: mathematical set and space with associated operators, enabling complex, multi-dimensional data analysis, often used in AI.

Vectorisation: the process of transforming unstructured data, such as text, into a numerical form, specifically a series of numbers forming a vector. This conversion is often carried out using embedding techniques, which are unsupervised learning algorithms that allow each element (such as a word in a text) to be represented by a point in a multi-dimensional vector space. This representation makes it possible to quantify the semantic relationships between elements and is easily exploited by AI.

Voice bot: a variant of the chatbot that interacts with users using voice commands.

Zero-shot mode: machine learning technique where a model can perform tasks for which it has received neither specific training nor a particular context.

Do you have any further questions?

Or if you would like to know more about how to create value from Artificial Intelligence, do not hesitate to mail us your query.

Digital Services

Digital Services is a business line within Orange Business, contributing to reliable and successful digital transformation for many organizations. Our joint mission is to help customers innovate and drive their business strategies in key digital domains, including Cloud, Customer Experience, Workspace, and Data & AI.
We assist them on their digital journey by providing advisory, end-to-end solutions, managed services, and professional services to ensure our customers' success.
We are digital natives, with innovation at the core of our business, which makes us a reliable partner close to our customers, leading them in their digital transformation challenges.

We support a wide range of industries in the private sector as well as the public sector. We have built a significant level of experience and understanding over the last 30 years in industries like Finance, Insurance, Life Sciences, Healthcare, Manufacturing, Travel & Transportation, Retail, and the Public Sector.
As always with Orange Business, our customers trust us for delivering end-to-end, sovereign, and sustainable solutions to turn their Operational Experience, Employee Experience and Customer Experience into business value.

Find out more how we can help you with your projects on:

digital.orange-business.com

Orange Business
Digital Services
Avenue du Bourget 3,
1140 Brussels
Belgium
marketing@bd-orange.com

Follow us on

orange™ Business