A SHORT GUIDE TO

ARTIFICIAL INTELLIGENCE

FOR BUSINESSES AND ENTREPRENEURS

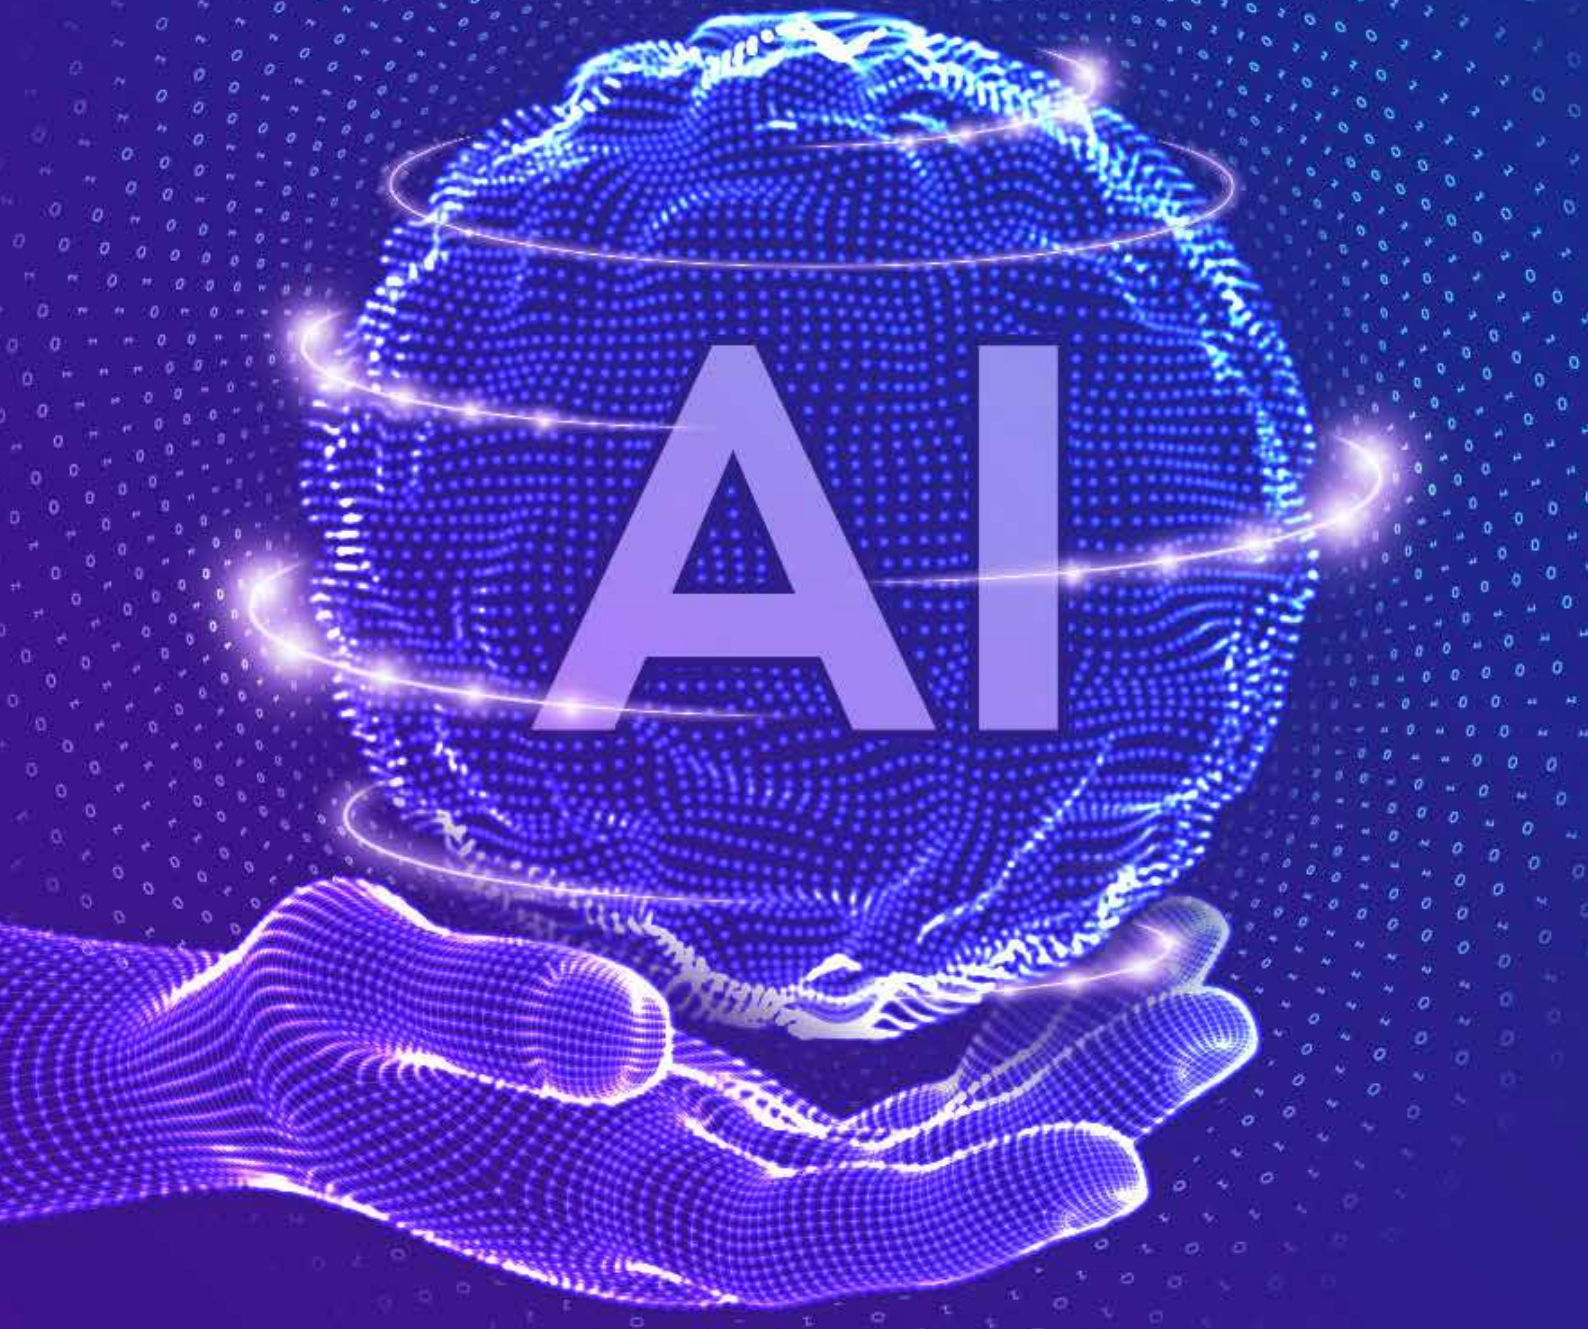

Vittorio Neri

Image by iuriimotov on Freepik

A SHORT GUIDE TO

ARTIFICIAL INTELLIGENCE

FOR BUSINESSES AND ENTREPRENEURS

Vittorio Neri

“ Non abbiate paura di immergere la punta del piede nell’oceano del domani. ”

Walt Disney

* immagine Wikipedia

INTRODUCTION

AI is evolving at an unprecedented pace, transforming technology and how we work, learn, and interact. Each week, new models emerge, research papers illuminate groundbreaking discoveries, and discussions grow more intense regarding AI's ethical, political, and societal implications. Meanwhile, the hardware that supports these advancements is striving to keep pace with increasingly sophisticated architectures designed to manage the growing complexity of AI models. As computational limits are stretched, new possibilities arise, paving the way for innovations that were once deemed impossible.

An Evolving Ecosystem.
The ecosystem surrounding AI encompasses the increasing energy demands required for AI model computation, efforts to mitigate the environmental impact of these technologies, and regulations designed to manage or prevent unethical or harmful applications of AI, such as the EU's AI Act or the UK's AI Regulation Bill. It also includes adoption methods for businesses, both large and small. Furthermore, we are witnessing how AI is influencing political strategies and geopolitical factors, with some nations leading the way while others strive to catch up.

The perfect moment to act.
This accelerated pace has shortened innovation cycles, which once allowed society more time to absorb new technologies. The sheer volume of information and techniques accompanying AI can feel overwhelming. Businesses, agencies, consultants, and organisations strive to understand how to translate these innovations into value while determining what AI could mean for their operations.

In marketing, my field of expertise, AI, clearly impacts numerous aspects, from data analysis to SEO, content creation, and strategy development. Similarly, other business functions, such as technical support, human resources, and sales, are transforming due to the deepening penetration of AI technologies. While this challenge may seem daunting, we can and must face it pragmatically—by understanding the basics and experimenting with one of the AI tools available today.

Perhaps I'm oversimplifying, but I firmly believe the Pareto principle applies perfectly here: understanding the 20% of concepts and tools that make the most significant difference will enable us to navigate the remaining 80% more efficiently. By doing so, we can establish a solid foundation that will assist us in contextualising and adopting future innovations, collaborating more effectively with suppliers and agencies, and maintaining the competitive edge that a well-informed use of AI affords.

Why do so many companies hesitate or ignore it?

Good question. Despite its immense potential, many companies still hesitate to adopt AI or even begin testing it. The reasons vary from the all-too-common excuse of "not having enough time" to genuine challenges rooted in a lack of knowledge, understanding, and direction. This leaves AI appearing too complex or distant from everyday reality. Additionally, some companies deliberately ignore AI, convinced it is irrelevant to their business needs or industry. They often perceive it as a tool solely suited to tech giants or cutting-edge sectors, far removed from the practicalities of their day-to-day operations.

There is a common belief among such businesses that AI offers little value for those not engaged in innovation-driven fields. For some, the constant pressures of daily operations overshadow the need to explore new technologies, leading to a mindset of "we don't have time for that." Others may dismiss AI as a passing trend or something that won't significantly impact their sector, preferring to focus on their established methods rather than investing in an unfamiliar technology.

However, ignoring AI entirely can have significant consequences. As technology becomes more pervasive and integrated into industries of all kinds, businesses that fail to engage with it risk falling behind their competitors. AI is transforming industries, creating efficiencies, reducing costs, and driving growth in traditional sectors. The idea that AI is optional or unnecessary could leave businesses unprepared to compete with those that have embraced it.

Addressing these barriers does not necessitate a large-scale transformation overnight. The first step often involves learning about AI's potential and experimenting with its applications in small, low-risk ways. The key is not to perceive AI as a disruptive challenge but as a practical tool for improving operations, enhancing creativity, and streamlining processes. While the reluctance to embrace change is understandable, inaction could ultimately prove more costly in an increasingly competitive landscape.

How this guide can help.

This short guide aims to outline, based on my experience, the essentials for exploring and assessing AI's potential. I've included simple yet practical examples to clarify concepts and illustrate how they can be applied. We'll also examine the challenges you might encounter and ways to address them. Although the journey may initially seem complex, adopting AI can become a rewarding and surprisingly enjoyable adventure with the right approach and an open mindset. One final note: this guide is not a technical or strategic manual on using AI. There is already extensive literature on the subject. The abundance of sources makes it necessary to have something that serves as a bridge between starting and gaining a deeper understanding. That's our aim today.

Happy reading!

01

AI.
A STRATEGIC ALLY FOR BUSINESSES.

Image by freepik

SOME MAY RECALL THE EMERGENCE OF THE FIRST COMPUTERS OR THE DAWN OF THE INTERNET. THESE MONUMENTAL MOMENTS MARKED A TURNING POINT, YET THEY WERE OFTEN ACCOMPANIED BY SIGNIFICANT INITIAL INVESTMENTS, ENCOURAGING MANY TO DELAY ADOPTION UNTIL COSTS BECAME MORE MANAGEABLE. FOR JUST A FEW EUROS A MONTH TODAY, HOWEVER, WE CAN ACCESS SOPHISTICATED YET USER-FRIENDLY ARTIFICIAL INTELLIGENCE TOOLS THAT SEAMLESSLY INTEGRATE INTO OUR DAILY ACTIVITIES. THESE TOOLS DRAW ON PAST EXPERIENCES TO MAKE EVERYTHING IMMEDIATELY ACCESSIBLE AND EASY TO USE. FOR MANY, EVERYTHING BEFORE NOVEMBER 2022, WHEN OPENAI INTRODUCED ITS AI MODEL TO CHAT WITH USERS, CHATGPT, ACCESSIBLE TO EVERYONE LIKE A STANDARD CHAT PLATFORM, REMAINS PRACTICALLY UNKNOWN.

BUT THAT DATE WASN'T A COINCIDENCE. SEVERAL FACTORS CONVERGED TO CREATE THAT PERFECT MOMENT: THE EVOLUTION OF COMPUTING POWER, THE VAST AVAILABILITY OF DATA FROM DIVERSE SOURCES (PUBLIC LIBRARIES, WEBSITES, BLOGS, COMMUNITIES AND INFORMATION FROM PUBLIC DATABASES LIKE PUBMED OR ARXIV) READY TO BE EFFECTIVELY UTILISED, ADVANCEMENTS IN SOFTWARE ARCHITECTURES (LIKE TRANSFORMERS, WHICH UNDERPIN LLMS), AND, NOT LEAST, THE INFLUX OF ABUNDANT FUNDING. THIS HAS MADE AI, AS WE KNOW IT TODAY, A REALITY.

AI, BORN IN THE 1950S, HAS EXPERIENCED HIGHS AND LOWS BUT HAS ALWAYS REACHED NEW MILESTONES. FROM TURING ONWARDS, ITS ADVENTUROUS HISTORY IS WORTH EXPLORING IN MORE DETAIL. TODAY, THANKS TO DATA AVAILABILITY, POWERFUL COMPUTERS EQUIPPED WITH SPECIALISED GPUS, AND ADVANCED ALGORITHMS LIKE DEEP NEURAL NETWORKS, WE CAN INTERACT WITH ARTIFICIAL INTELLIGENCE TOOLS AS EFFORTLESSLY AS WE CHAT ON WHATSAPP.

IF YOU'VE NEVER USED AN AI PLATFORM, IMAGINE IT AS A CHAT INTERFACE: A MESSAGING SOFTWARE THAT CAN RESPOND TO ALMOST ANY OF YOUR REQUESTS. THIS MIGHT SOUND OVERLY SIMPLISTIC FOR THOSE WHO HAVEN'T TRIED TOOLS LIKE CHATGPT, CLAUDE, OR GEMINI, BUT IT'S NOT FAR FROM THE TRUTH. YOU TYPE IN YOUR LANGUAGE (KNOWN AS NATURAL LANGUAGE), AND ON THE OTHER SIDE, YOU HAVE A POTENT TOOL THAT ENABLES YOU TO:

1. AUTOMATE TIME-CONSUMING TASKS, SUCH AS REVIEWING FINANCIAL STATEMENTS OR REPORTS, FREEING UP VALUABLE RESOURCES FOR MORE STRATEGIC ACTIVITIES.

2. BOOST CREATIVITY BY OFFERING IDEAS, PERSPECTIVES, AND SOLUTIONS YOU MIGHT NOT HAVE THOUGHT OF, OFTEN EXCEEDING EXPECTATIONS.

3. SPEED UP CONTENT CREATION OR THE ANALYSIS OF COMPETITOR DOCUMENTS.

WHETHER CREATING ADVERTISING CAMPAIGNS, ANALYSING BUSINESS DATA, OR PERSONALISING CUSTOMER EXPERIENCES, AI PROVIDES PRACTICAL TOOLS TO PRODUCE READY-TO-USE CONTENT OR SERVE AS A FOUNDATION TO REFINE OUR WORK FURTHER.
IT'S ALL TOO FASCINATING TO AFFORD THE LUXURY OF IGNORING IT.

Debunking misconceptions while staying realistic.

Let's address some common misconceptions about artificial intelligence, often perpetuated by mainstream media, where AI is frequently portrayed as:

- A threat to jobs, with fears that it could replace humans.
- An out-of-control tool, acting without supervision.
- A technology reserved for technical experts only.

Starting with the last point. For many platforms, artificial intelligence is more accessible than one might think. Thanks to the adoption of natural language interfaces that allow users to interact with AI tools using everyday language rather than code, along with the introduction of chat-like interfaces, the complexity of using these systems has significantly decreased. A glance at AI's widespread adoption today suffices to confirm this. Additionally, countless courses have emerged to assist individuals in understanding how to use the most common platforms, further lowering the barrier to entry.

Regarding the notion that AI operates autonomously, it's undeniable that AI will evolve in this direction, particularly with tools known as agents and, eventually, robots equipped with more advanced AI. These robots, powered by increasingly sophisticated intelligence, could become indispensable in a world facing growing shortages of skilled and unskilled labour. In sectors like healthcare, manufacturing, and logistics, where the demand for workers often outstrips supply, robots with enhanced AI capabilities could bridge the gap, performing repetitive, labour-intensive, or even hazardous tasks for humans. This development has the potential to alleviate some of the workforce challenges that many industries currently face, ensuring businesses can continue to grow and serve their markets effectively.

Agents, which are essentially pre-programmed AI systems designed to carry out tasks autonomously, are already on the rise. Whether working independently or collaborating with other agents, they promise to significantly enhance productivity by automating complex workflows. While still in their infancy, these tools are rapidly advancing, fuelled by substantial investment from major tech companies. It's not far-fetched to envision a future where an "app store" model emerges for AI, allowing businesses and individuals to purchase tailored agents or even entire AI systems designed to meet specific needs, much like we currently buy software or cloud services. This commodification of AI could dramatically lower entry barriers, enabling smaller businesses and even individuals to access and leverage AI technology in previously unimaginable ways.

Moreover, the broader evolution of AI and its integration into robots and agents will likely transform how we think about and use technology in our lives. As AI becomes increasingly efficient, adaptable, and customisable, it's possible that owning a personal AI assistant or a task-specific robot will be as commonplace as owning a smartphone or subscribing to a streaming service. These developments will reinforce AI's role as a practical and accessible tool rather than a distant or intimidating technology.

For the time being, human supervision remains critical, not only to ensure the effective use of these technologies but also to validate their outputs and maintain ethical standards. As we've explored, AI still requires human oversight, particularly when validating the content and decisions generated by these systems. This partnership between human expertise and AI efficiency is key to ensuring that AI can deliver real value while operating within safe and ethical boundaries.

AI does not replace human intelligence but enhances it when utilised effectively. While it serves as a valuable tool for increasing productivity, the validation of results remains firmly in the hands of knowledgeable individuals who oversee the output and ensure its quality. What AI does replace are specific tasks—often repetitive, low-value, and automatable. This is why I strongly advocate the adoption of AI sooner rather than later, to avoid becoming passive consumers of the technology and risking a competitive disadvantage against those who have acquired even a basic proficiency in its use.

As for AI's impact on jobs, while there are estimates and projections, it's still too early to draw concrete conclusions. Typically, where automation takes hold, roles and skillsets undergo revision. This will likely happen sooner rather than later as AI becomes an integral part of daily work and increasingly pervasive in society. Once again, understanding and mastering AI should be prioritised in our work, personal, and professional lives.

02

UNDERSTANDING ARTIFICIAL INTELLIGENCE.

Image by ojosujono96 on Freepik

AT THE OUTSET, WE NOTED THAT KNOWING 20% OF THE INFORMATION ENABLES YOU TO HANDLE 80% OF ACTIVITIES. UNDERSTANDING THE ORIGINS OF AI IS PART OF THAT ESSENTIAL 20%. THEREFORE, LET'S FAMILIARISE OURSELVES WITH THE "PLAYING FIELD" AND GAIN CLARITY ON THE TERMINOLOGY AND ITS MEANINGS TO BETTER CONTEXTUALISE THE VARIOUS CONCEPTS.

ARTIFICIAL INTELLIGENCE (AI) CAN BE DEFINED IN VARIOUS WAYS, YET IT FUNDAMENTALLY REPRESENTS A BRANCH OF COMPUTER SCIENCE THAT DEVELOPS SYSTEMS CAPABLE OF PERFORMING TASKS TYPICALLY REQUIRING HUMAN INTELLIGENCE. SUCH TASKS MAY INCLUDE RECOGNISING OR GENERATING IMAGES, ANALYSING DOCUMENTS, COMPREHENDING TEXT, AND MAKING COMPLEX DATA-DRIVEN DECISIONS. FURTHERMORE, AI SIGNIFIES A NOTABLE ADVANCEMENT IN CREATING MACHINES THAT CAN BEHAVE APPROPRIATELY IN DIVERSE AND DYNAMIC SITUATIONS, ADAPTING TO CONTEXT AND RESPONDING INTELLIGENTLY TO NEW CHALLENGES. THIS ALIGNS WITH THE PERSPECTIVE OF DEMIS HASSABIS, CEO OF DEEPMIND, WHO DESCRIBES AI AS "THE SCIENCE OF CREATING INTELLIGENT MACHINES." THIS DEFINITION HIGHLIGHTS AI'S BROADER AMBITION—NOT MERELY TO REPLICATE HUMAN TASKS BUT TO DEVELOP SYSTEMS THAT DEMONSTRATE GENUINE PROBLEM-SOLVING CAPABILITIES, DECISION-MAKING ADAPTABILITY, AND LEARNING POTENTIAL ACROSS DIVERSE ENVIRONMENTS.

WE CAN DISTINGUISH TWO BROAD CATEGORIES OF AI:
THE FIRST IS **GENERAL AI**, ALSO KNOWN AS AGI (ARTIFICIAL GENERAL INTELLIGENCE), WHICH REFERS TO SYSTEMS CAPABLE OF PERFORMING ANY INTELLECTUAL TASK A HUMAN CAN DO WITHOUT HUMAN INTERVENTION. ALTHOUGH THIS CONCEPT IS STILL THEORETICAL TODAY, MANY EXPERTS CONSIDER IT A FUTURE MILESTONE WITH THE POTENTIAL TO CREATE FULLY AUTONOMOUS AND INDEPENDENT AI SERVICES.

AS PREVIOUSLY MENTIONED, SOME PLATFORMS ARE BEGINNING TO INTEGRATE ASSISTANTS CALLED AGENTS, DESIGNED TO COORDINATE TASKS WITH OTHER AGENTS. THESE IMPLEMENTATIONS BRING US CLOSER TO AGI. RECENTLY, THE TERM **ASI (ARTIFICIAL SUPER INTELLIGENCE)** WAS INTRODUCED TO DESCRIBE SYSTEMS THAT WOULD SURPASS HUMAN INTELLIGENCE BY A SIGNIFICANT MARGIN IN TERMS OF POWER AND INDEPENDENCE. HOWEVER, AS OF TODAY, SUCH SYSTEMS DO NOT EXIST. THEIR DEVELOPMENT TRAJECTORY WILL INEVITABLY BE ACCOMPANIED BY REGULATIONS AND CONTROL STRUCTURES TO PREVENT UNMANAGEABLE OUTCOMES.

THE SECOND CATEGORY IS **SPECIALISED AI**, ALSO KNOWN AS NARROW AI OR WEAK AI. THIS ENCOMPASSES TOOLS AND SYSTEMS DESIGNED FOR SPECIFIC TASKS, SUCH AS NATURAL LANGUAGE PROCESSING, EXEMPLIFIED BY CHATGPT; VOICE RECOGNITION, UTILISED BY VIRTUAL ASSISTANTS LIKE ALEXA AND SIRI; IMAGE GENERATION, ENABLED BY TOOLS LIKE DALL-E AND MIDJOURNEY; AND PREDICTIVE ANALYTICS, WHICH DRIVE SOCIAL MEDIA FEED SUGGESTIONS OR MOVIE RECOMMENDATIONS ON STREAMING PLATFORMS. EVEN SPAM FILTERS THAT IDENTIFY AND BLOCK UNWANTED EMAILS FIT INTO THIS CATEGORY.

WE USE SPECIALISED AI IN OUR DAILY LIVES AND BUSINESSES. IT IS BUILT ON VARIOUS TECHNOLOGIES, WHICH WE WILL EXPLORE IN THE FOLLOWING SECTIONS.

GIVEN THE RAPID EVOLUTION OF AI, IT'S IMPORTANT TO NOTE THAT THESE TERMINOLOGIES SHOULD ALWAYS BE APPROACHED WITH SOME FLEXIBILITY. PRODUCERS AND INDUSTRY COMMENTATORS OFTEN INTERPRET AND DESCRIBE THESE CONCEPTS DIFFERENTLY BASED ON THEIR PERSPECTIVES OR ROLES.

Machine Learning, Deep Learning, and Generative Artificial Intelligence.

To understand how Artificial Intelligence (AI) works, imagine a hierarchical structure, like a Matryoshka doll: each layer contains a more specific technology. This layered simplification helps us grasp how each part of AI contributes to the broader technological ecosystem and how each subcategory is designed to solve particular problems and address increasingly complex challenges. Our ultimate goal is to reach the type of AI that interests us most: **Generative AI.**

INTELLIGENZA ARTIFICIALE

Tecnologie che permettono alle macchine di imitare l'intelligenza umana.

MACHINE LEARNING

Le tecnologie imparano dai dati per migliorare le prestazioni.

DEEP LEARNING

Forma avanzata di Machine Learning che usa reti neurali (operazioni matematiche) per analizzare dati complessi.

IA GENERATIVA

IA che può creare contenuti come testi e immagini, in base a quello che ha appreso dai dati.

LLM

Modelli di IA che lavorano su testo o multimodali.

DIFFUSION MODELS

Modelli di IA che lavorano anche su immagini.

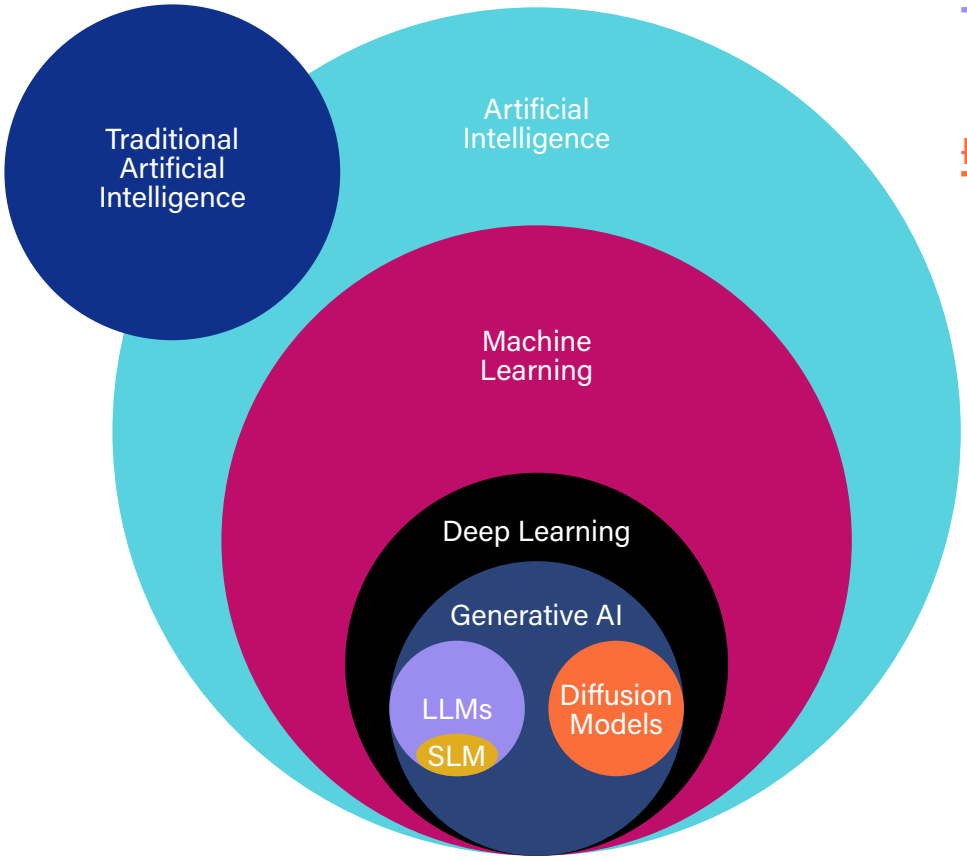

Artificial Intelligence. The Foundation.

As previously mentioned, AI is a branch of computer science that develops technologies to simulate human abilities like reasoning, learning, and problem-solving. It is the broadest category under which various approaches and techniques are classified. AI's goals range from solving simple tasks to tackling complex challenges, such as medical diagnoses or autonomous driving. This broad spectrum makes AI a rich and diverse field with applications spanning industries, medicine, marketing, finance, and entertainment.

Machine Learning, AI That Learns from Data.

Within AI, we find Machine Learning (ML), also known as automatic learning—a subfield that enables machines to learn from data without being explicitly programmed for every task. Machine Learning leverages neural networks and algorithms (think of the networks as the structure of a house and the algorithms as the instructions for building and operating it) to mimic the workings of the human brain, identifying patterns and relationships in data, such as categorising data into groups or similar characteristics. Typical applications include spam filters, recommendation engines on websites or streaming platforms, image recognition, and disease diagnosis based on thousands of clinical data points used during training.

Machine Learning is significant because, until a few years ago, AI systems were built on fixed rules defined by experts (known as expert systems, where the system relied on domain-specific expertise). A well-known example is Eliza, one of the earliest medical chatbots, which responded using conditional structures—essentially, "If this happens, do that."

However, this approach had its limitations, as every change or update necessitated adjustments to the code. With increasing complexity and growing data volumes, alternative methods became crucial. Machine Learning allows systems to adapt autonomously to new data and to update themselves as they gather more information, rendering them more flexible and efficient.

Rule-based AI—commonly called symbolic AI—has recently been re-evaluated as a complementary method to machine learning. This hybrid approach aims to address the limitations of each by combining their strengths, particularly in areas such as logic, deductive reasoning, and the transparency of AI model reasoning (often termed "black box" systems due to the obscurity surrounding the internal processes within AI systems). It remains to be seen how this will evolve.

Deep Learning. The Revolution.

Deep Learning (DL) is a subset of Machine Learning that employs deep neural networks. These networks can process vast amounts of data to identify complex patterns and tackle more intricate tasks. Unlike traditional Machine Learning, DL becomes essential when substantial computational power is required to manage considerable quantities of data.

In contrast to simpler neural networks, deep networks consist of multiple layers (hence the term "deep"), each responsible for a different level of analysis—from basic details to more abstract concepts.

Each layer of the neural network plays a specific role. The initial layers (input layers) identify basic features, such as lines, colours, or text patterns. Subsequent layers (hidden layers) process more complex combinations, eventually recognising entire objects or situations and providing the final output (output layer).

Deep Learning has enabled computers to perform highly complex tasks (which come naturally to humans), such as facial recognition, automatic language translation, emotion analysis (sentiment), and trend prediction on social media. To achieve this, systems undergo training to adjust their parameters to provide correct responses. Imagine the process as a gigantic mixer optimising billions of parameters to produce the final output. These parameters act like tiny "switches" fine-tuned during the learning process to improve the model's accuracy. The more data the system receives, the better the neural network becomes at making accurate predictions or recognising intricate patterns.

Consider facial recognition, as employed in our phones and surveillance systems. These algorithms analyse millions of images during training to comprehend the distinct features of human faces. Through ongoing enhancement, they can identify individuals with remarkable accuracy.

Other applications encompass autonomous vehicles, which utilise deep neural networks to recognise traffic signs, pedestrians, and other vehicles. These systems analyse substantial amounts of data from various sensors and make split-second decisions to ensure safety.

Generative AI: From Understanding to Creation.

Deep Learning has also paved the way for **Generative AI** systems capable of producing new content such as text, images, videos, and music. This has undoubtedly been a stroke of genius, achieved by combining deep neural networks with innovative algorithms (such as the Transformer utilised by ChatGPT) and training these systems on text, providing them with millions of examples and solutions to assist them in recognising patterns and semantic structures.

The Transformer, developed by Google, introduced the attention mechanism, enabling an AI model to assign different weights to elements in a text based on their relevance to the context. This mechanism revolutionised text processing by allowing a more accurate understanding of the relationships between words and phrases. Initially focused on translation tasks, it utilised algorithms to significantly enhance machine translation quality. However, an alternative application of the Transformer and the attention mechanism, as adopted by OpenAI, the creator of ChatGPT, revealed broader potential. By training models on vast amounts of data, researchers discovered that these systems could transcend translation and demonstrate reasoning and abstraction capabilities, making them suitable for complex linguistic and cognitive tasks. This insight laid the groundwork for developing large language models (LLMs), fundamentally altering how AI interacts with and processes human language.

Examples of Generative AI today include ChatGPT, Gemini (for text and more), and DALL-E, Flux, or Midjourney (for images). These models process vast amounts of data, applying learned patterns to generate new content now that computational power and data availability are no longer bottlenecks. Another breakthrough was the decision to mimic the workings of the human brain using statistical methods rather than attempting to create anthropomorphic models, as was previously done with the initial versions of robots (which we can recall from many films).

As we said, the most famous generative models today are **Large Language Models (LLMs)**, which specialise in text processing and are now multimodal, meaning they can work with images, audio, and code both in input and output. (ChatGPT is a multimodal model because it handles text, images, and code.) Other examples include **Diffusion Models**, used for creating images and videos. These models start with a random noise representation and generate detailed and realistic images through a series of refinements. Imagine creating illustrations of entirely fictional landscapes that look like photographs of real places.

This also highlights the distinction between Generative AI and traditional software. Take a text editor like Word, for instance: currently, you cannot ask it to generate new text, which is precisely what an LLM can do. This is why many conventional software programs will soon be equipped with generative features. For example, the Microsoft suite includes Copilot, which provides AI capabilities in each Microsoft application. Google is also releasing Gemini on its applications.

Given the above, it's clear how Generative AI impacts numerous sectors, from the entertainment industry, where it creates special effects and virtual sets, to publishing, where it writes articles or summarises news stories.

Moreover, Generative AI is entering education, forcing a rethinking of teaching methods regarding using such tools for learning and understanding what they produce. This presents a significant challenge for education and schools in general.

Predictive vs Generative Artificial Intelligence.

One final note: Artificial Intelligence has become a magical buzzword, often accompanying many software products on the market. However, if you recall the chapter where we discussed AI's structure and various components (Conditional, Machine Learning, Deep Learning), you'll see it has been with us for quite some time.

We didn't have Generative AI—not as we know it today—but Machine Learning and Deep Learning have long enabled predictive artificial intelligence. Predictive AI appears in dashboards, data visualisations, market trend analyses, stock management, lead scoring, web traffic insights, or sales forecasts. These insights are primarily powered by predictive AI systems, which use statistical algorithms and Machine Learning models to provide answers based on mathematical patterns rather than creating new content.

Often, these models are embedded in our software, making them invisible to us, but they are there. Some examples of predictive AI in use include:

- Forecasting product demand to optimise stock levels and reduce waste.
- Analysing historical sales data to identify seasonal trends.
- Detecting unusual behaviour for financial fraud prevention.
- Monitoring production lines to identify defective products.
- Calculating customer churn risk through retention analysis.

In summary, predictive AI helps businesses predict the future by analysing past data, while generative AI enables the creation of something entirely new. These are two very different tasks, but they both enhance productivity and competitiveness and are often integrated to create even more effective solutions.

03

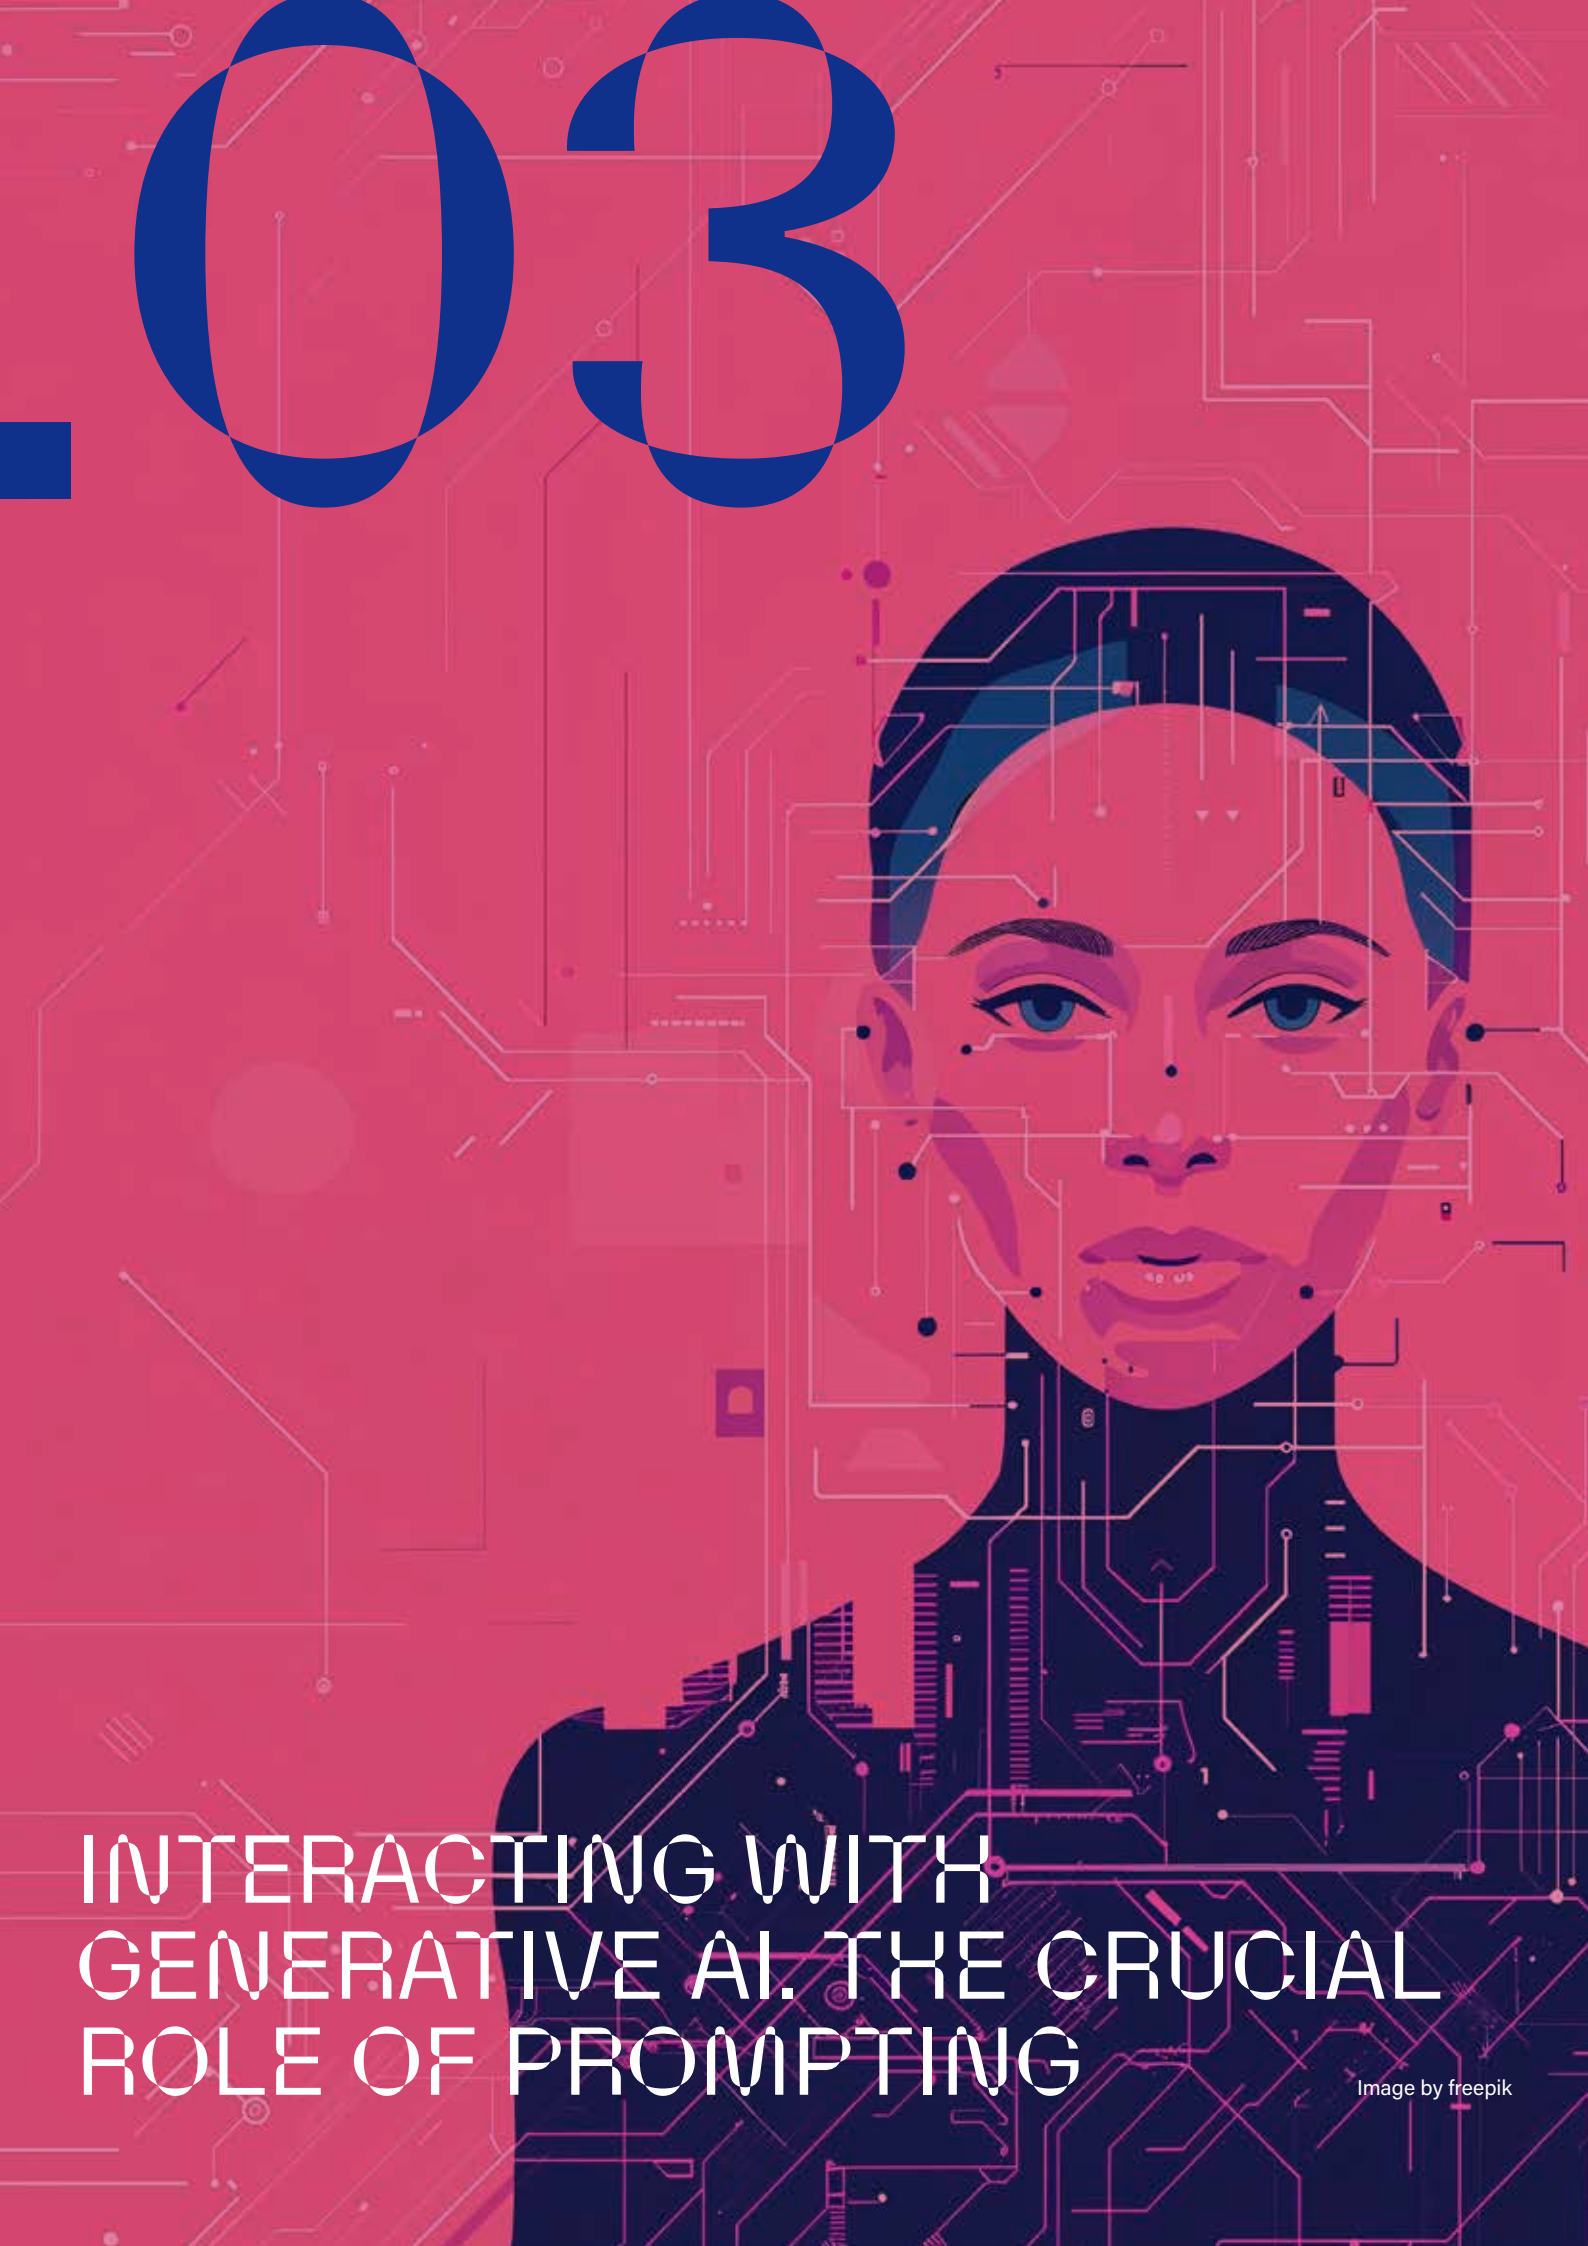

INTERACTING WITH GENERATIVE AI: THE CRUCIAL ROLE OF PROMPTING

Image by freepik

ANOTHER KEY ASPECT OF THE 20% YOU MUST UNDERSTAND IS COMMUNICATING EFFECTIVELY WITH GENERATIVE AI MODELS.

UNLIKE TRADITIONAL SYSTEMS, GENERATIVE AI MODELS AREN'T PROGRAMMED IN THE CONVENTIONAL SENSE; INSTEAD, INTERACTION OCCURS THROUGH PROMPTING. THIS TECHNIQUE PROVIDES CLEAR AND STRUCTURED INSTRUCTIONS TO ACHIEVE THE MOST ACCURATE RESPONSES POSSIBLE.

SPEAKING TO AN LLM (LARGE LANGUAGE MODEL) IS A NEW FORM OF "PROGRAMMING": THE MORE PRECISE THE PROMPT, THE BETTER THE RESULT. A WELL-CRAFTED PROMPT SIGNIFICANTLY AFFECTS THE QUALITY OF THE GENERATED CONTENT, OFFERING DETAILED GUIDANCE THAT SHAPES THE MODEL'S "CREATIVITY."

THE WORD "PROMPT" COMES FROM THE LATIN "PROMPTUS," MEANING "READY," AND IT ENABLES US TO PREPARE THE MODEL TO RESPOND OR SUGGEST HOW TO STRUCTURE ITS ANSWERS.

The Probabilistic Nature of Responses.

One fundamental aspect to understand is that LLMs operate probabilistically or statistically. As previously mentioned, adopting this approach was a stroke of genius. These models don't provide an exact answer but the most probable one, though the two often align. The text generated by an LLM relies on the relationships between words and concepts, predicting the most suitable one based on the context.

For example, if we write, "The dog runs in the...", the model might select "field" as the most probable solution. However, if we provide more details specifying that the dog is in an urban setting, the model might instead choose "park," making the response more appropriate and contextual.

This is fascinating because the model learns the probability of a correct answer through the vast quantities of text processed during training. The quality of the prompt further refines this. The model's data streams for generating responses encompass a wide range of information, including books, projects, strategies, biographies, patents, research papers, and more. This extensive collection renders LLMs a "repository of human knowledge"—at least regarding what is available online (as far as we know).

For this reason, a good prompt must nudge the model to respond accurately and personally, avoiding generic or fabricated answers.

How Words Become Answers.

Building on the earlier example of the dog running in the field or park, it's fascinating to understand how the model arrives at these responses. It begins with the prompt, broken down into smaller units called tokens. These tokens are the building blocks for the model: they can be whole words, parts of words, or symbols. For instance, the phrase "The dog runs in the..." is tokenised into: ["The", "dog", "runs", "in"]. For more complex terms, tokenisation can go even further, breaking a word like "running" into ["run-", "-ning"].

After tokenisation, the model transforms each token into an embedding, a numerical representation that places the words in a multidimensional space. In this space, words with similar meanings (like "dog" and "cat") are positioned close together, while less related words (like "dog" and "aeroplane") are farther apart. This step enables the model to "understand" context and semantic relationships.

Next, the Transformer takes over, analysing the context provided by the prompt. Here, an iterative process unfolds: each token is compared to the preceding ones, allowing the model to calculate which word or phrase makes the most sense based on probabilities. This continuous processing flow ensures the generated response aligns with the context and available information.

The final step is the Output Layer, which predicts the most probable next token. This is how the model completes the phrase: "The dog runs in the..." becomes "field" if the context is neutral or "park" if the context suggests an urban setting.

This process—from the prompt through tokenisation, embeddings, and internal iterations—illustrates how crucial it is to provide a clear and detailed prompt. As demonstrated in the earlier example, the model can only produce precise and contextualised responses.

The Context Window and the AI's Processing Range.

When using an artificial intelligence model like an LLM (Large Language Model), such as ChatGPT or Claude, there's a limit to the amount of information the model can process at any given time. This limit is the context window and represents the model's capacity to "remember" and work with the data provided in a single interaction.

Think of having a conversation with a colleague: you can discuss a topic in great detail, but if too much information is shared simultaneously, it becomes harder to keep track of everything. The context window works in much the same way. For an LLM, everything within the window is accessible for generating a response, while anything "outside" of it is effectively forgotten.

The context window defines the maximum number of tokens (words, parts of words, or symbols) the model can process. For example, a model with a 4,000-token context window can handle roughly 3,000 words (since tokens often represent parts of words). Advanced models, however, can process up to 32,000 tokens or more, allowing them to handle long documents, such as reports or contracts, within a single interaction.

Why is this important? Understanding how the context window works is key to getting high-quality responses. Providing relevant information, structuring text by starting with the most critical details, and avoiding information overload helps the model focus on the essential points.

If the number of tokens exceeds the context window's limit, the model discards the earliest parts of the conversation to make room for newer input—a process called truncation. When working with long texts, it's helpful to break them into manageable sections or to summarise key points as you go.

For instance, if you want to write a detailed product description and start with 2,000 words of context, you might use the available window before getting the desired output. A better strategy would be to provide a concise introduction, followed by relevant details, ensuring the model isn't overwhelmed with too much information at the outset.

Understanding and working within the boundaries of the context window can help you make the most of AI tools and ensure accurate, high-quality responses. However, this aspect constantly changes, and new models have increasingly large context windows. To expand the content windows and optimise computational efficiency, new architectures incorporate interconnected modules for short-term, long-term, and working memory, enabling it to handle larger contexts more effectively. A standout feature is its dynamic long-term memory, which adapts by selectively storing new information based on its "surprise" factor—how distinct it is from existing knowledge. An adaptive forgetting mechanism optimises memory use by removing outdated or irrelevant data.

This innovation moves AI closer to mimicking human-like reasoning, improving its capacity to efficiently process complex, context-rich information.

The "Hallucinations" of AI Models.

LLM models are remarkably generous machines. They consistently provide an answer, even when they lack the required information. While this can be beneficial in many scenarios, it carries a significant risk: the possibility of the model "inventing" answers, leading to what are known as "hallucinations." These responses seem credible—written in a fluid and logical manner—but are entirely false or unfounded.

For example, asking a model to cite a document or a specific source it has never encountered might generate a plausible but entirely fabricated response. It could even include details like article titles or author names that seem authentic. This happens because the model does not "know" like a human does; instead, it generates responses based on probabilistic correlations.

A well-structured prompt can help mitigate the risk of hallucinations by steering the model towards accurate and reliable responses. For instance:

- Ask the model only to cite information it is "confident" about.
- Explicitly request sources, using prompts like, "Include sources if available."
- Provide context or specific details, avoiding vague questions that might encourage the model to improvise.

However, even with a carefully crafted prompt, the risk of errors cannot be eliminated entirely. This is why it is crucial always to verify responses when dealing with critical decisions, sensitive data, or content intended for publication. AI is a powerful tool, but human validation remains essential to ensure the quality and reliability of its outputs.

Conversational vs Reasoning LLMs: What's the Difference?.

As AI evolves, we're starting to see a distinction between different types of Large Language Models (LLMs). While early models were primarily built for conversation, today, we have two main categories: **conversational** and **reasoning models**.

Conversational LLMs—much like the early iterations of AI chat assistants—are crafted to engage in natural conversations. Their main aim is to facilitate smooth dialogue, respond promptly, and adapt to various tones and styles. These models excel at sustaining engaging and dynamic discussions but may not explore complex reasoning in depth.

Reasoning LLMs, on the other hand, take things a step further. Instead of just answering based on context, they **analyse problems, break down complex tasks**, and provide more structured, step-by-step responses. This makes them better at logical thinking, coding, data analysis, and problem-solving—tasks requiring more than quick responses.

The best part? These two approaches are starting to merge. AI models today are becoming conversational and capable of deeper reasoning, meaning future AI assistants won't just chat with you—they'll help you think, plan, and make better decisions.

Image Prompting.

Prompting isn't limited to generating text—it also plays a pivotal role in creating images. Models like DALL-E and MidJourney rely on the details provided in a prompt to craft specific outputs, such as colours, artistic styles, objects, and particular settings.

For instance, a prompt like "a cat sleeping on a red baroque sofa" will produce a vastly different result from "a cat playing on a sunny field of grass." This ability to translate textual descriptions into visual creations highlights why mastering prompting is essential for anyone looking to harness generative AI effectively.

The Art of Prompting: Effectively Communicating with AI.

Clearly, the effectiveness of a generative AI system largely depends on the quality of the instructions it receives. A well-crafted prompt isn't just a simple question (although that's sometimes enough); it's a strategic way of communicating with the AI to guide its responses. Crafting a strong prompt can turn a mediocre result into an exceptional one.

For instance, a generic prompt like "Write an article on marketing" offers a basic starting point. However, a more detailed prompt, such as "Write a 500-word article on digital marketing for small businesses, including practical examples and actionable tips," produces far more focused and relevant content.

Key Elements of an Effective Prompt.

1. **Specificity**- Provide precise details about the topic, format, and desired outcome.

2. **Style and Format** - Specify the tone (formal, creative, technical) and format (list, paragraphs, chart).

3. **Audience** - To assist the AI in adjusting the tone and complexity, specify the target audience, such as "small business owners" or "university students."

4. **Iterative Flip-Flop** -Don't settle for the first attempt. Collaborate with the AI to refine the prompt by asking, "The response isn't complete—can you add an example?" or "Rephrase this to be more persuasive." This iterative process enhances results progressively.

5. **Context and Constraints** - Provide relevant data, references, or limits, such as "Focus on 2023 trends" or "Keep it under 200 words."

Mastering these elements improves the output quality and allows users to unlock the full potential of AI systems in their workflows.

Learning Prompting.

Mastering the art of prompting is essential for maximising the potential of AI systems. While expertise isn't necessary, dedicating time to improve this skill can produce remarkable results. The good news is that a variety of free and paid courses are available to help you learn the fundamentals of prompting. These courses instruct on how to construct better prompts and offer advanced techniques for unlocking the full capabilities of generative AI.

Crafting effective prompts is like learning how to communicate with an exceptionally brilliant but slightly reserved colleague who needs clear direction to shine. Investing in prompting practice means fewer iterations, higher-quality results, and the creation of high-value content or solutions. It's not a complicated task but rather a strategic skill that turns AI from a mere tool into a faithful operational ally.

Like any form of communication, prompting is an art that improves with practice. The great thing is that it's a skill you can pick up quickly, thanks to available resources and, most importantly, by directly interacting with AI to craft progressively better prompts.

04

PRACTICAL APPLICATIONS AND BENEFITS OF GENERATIVE AI.

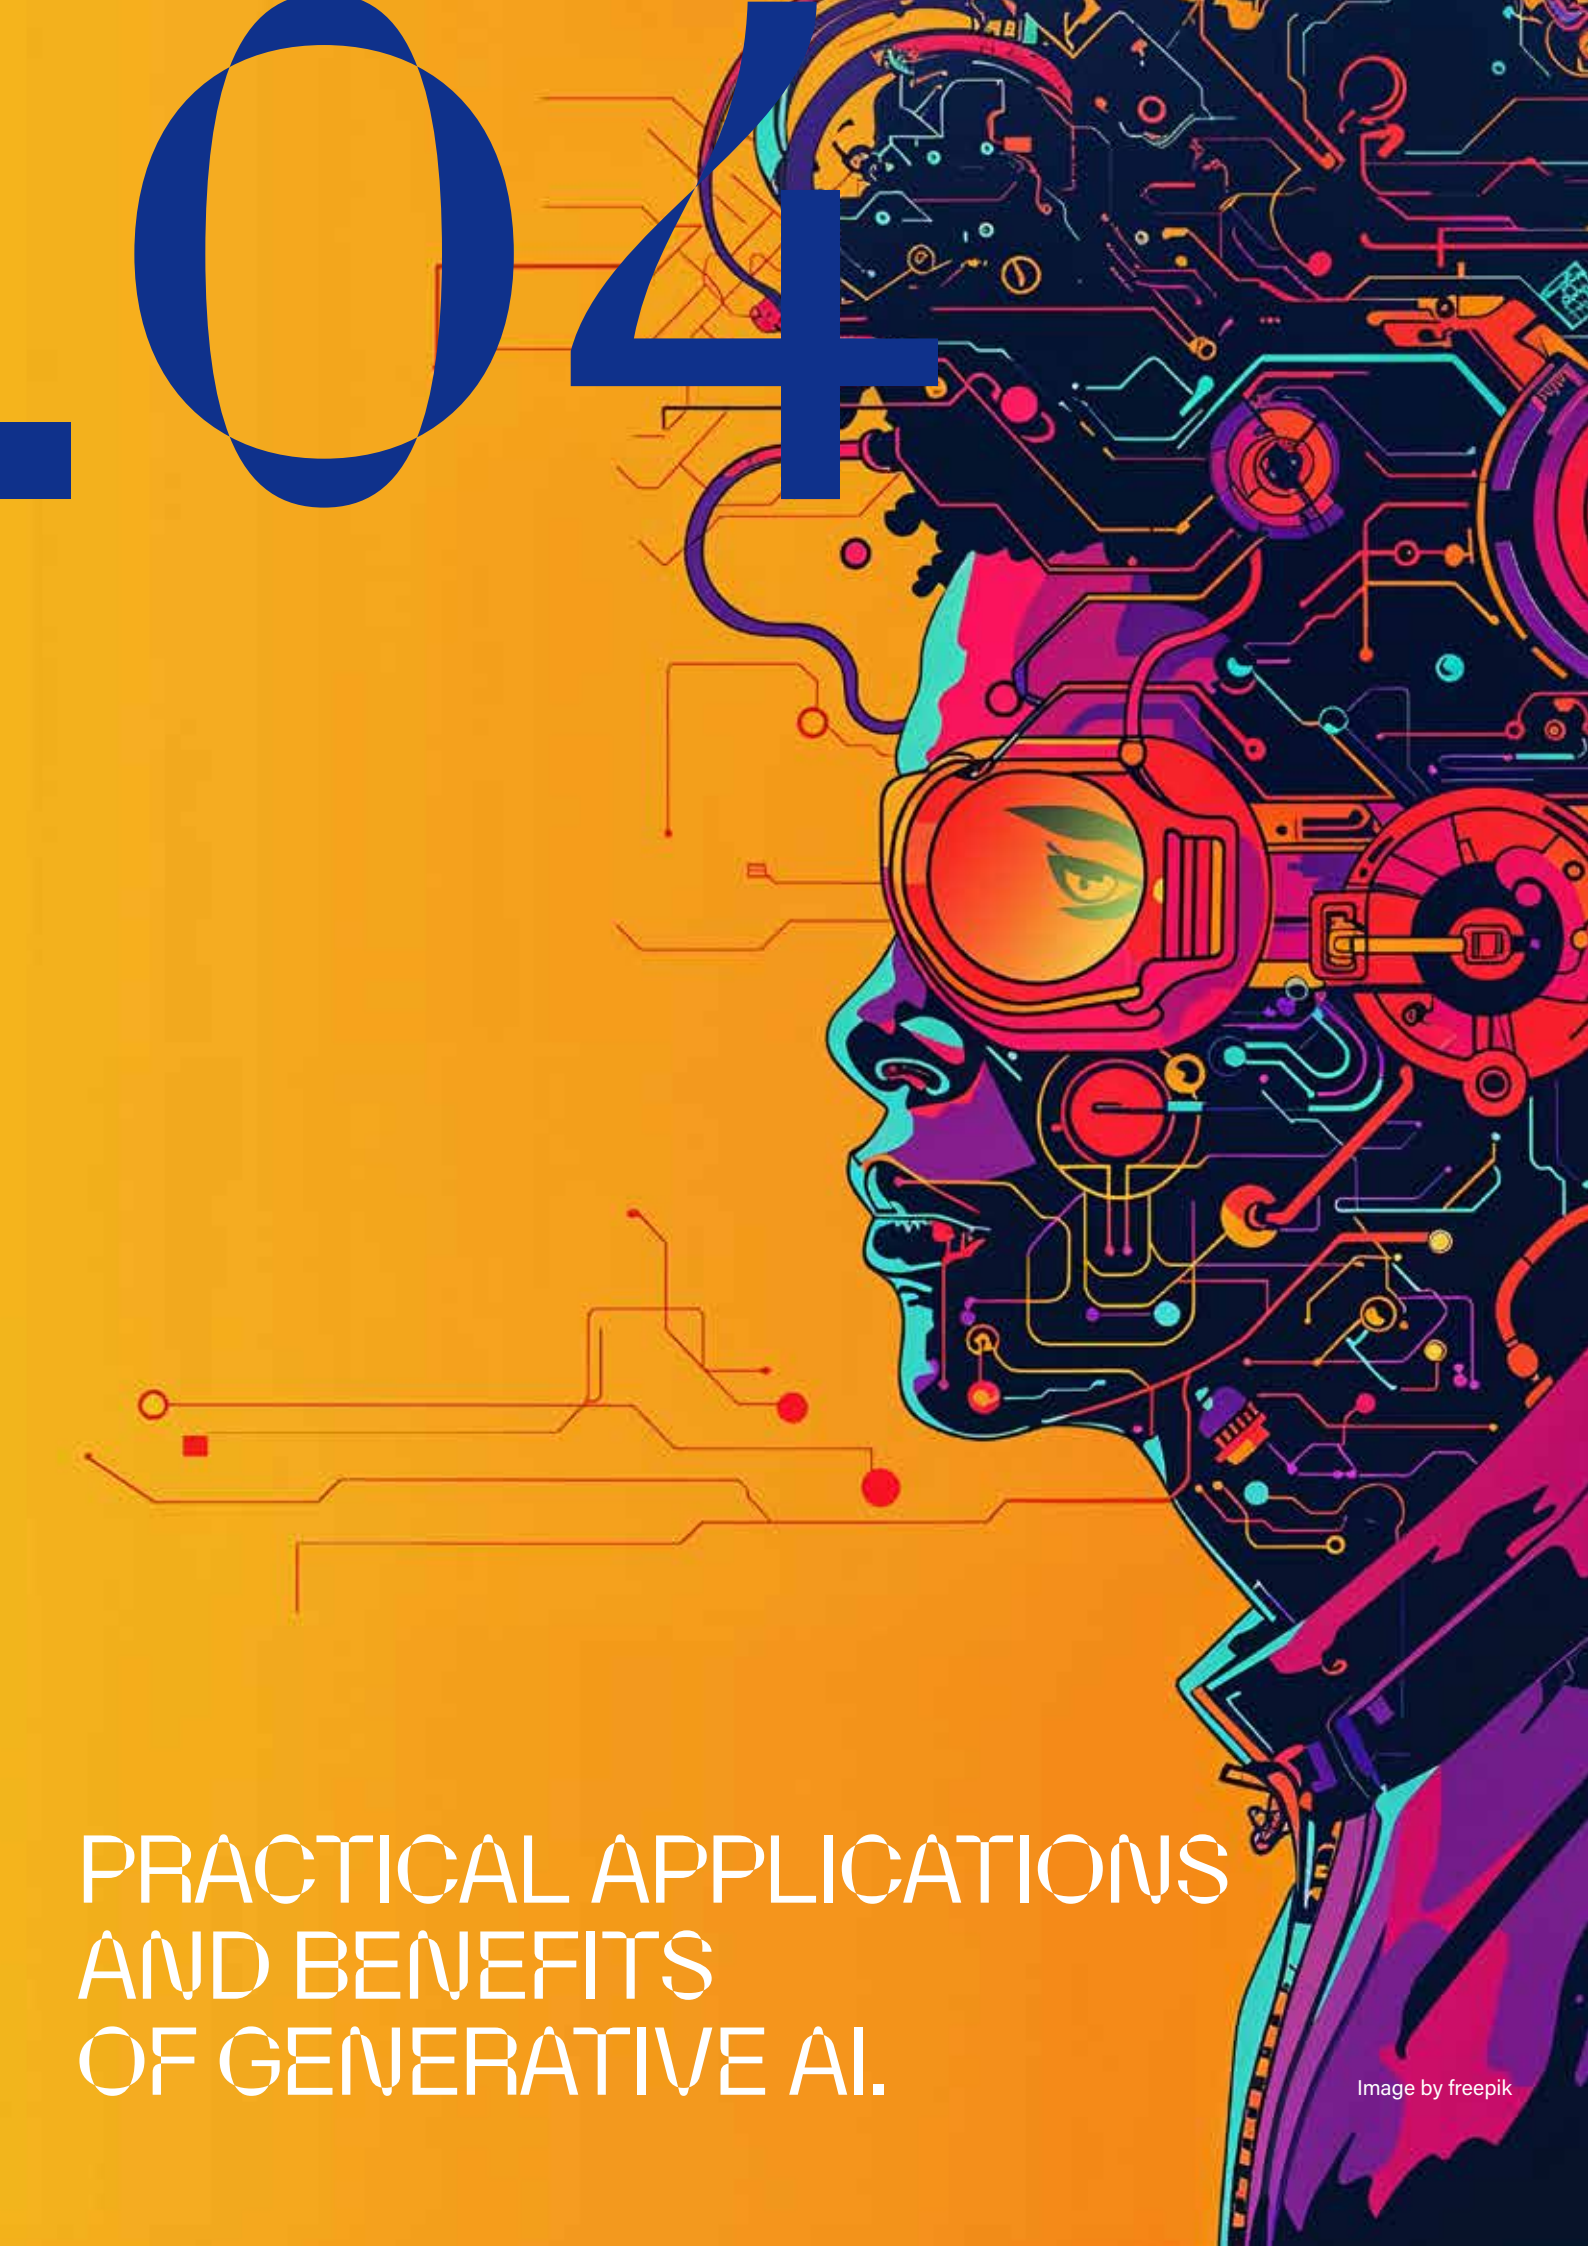

Image by freepik

LET'S DIVE INTO WHAT GENERATIVE ARTIFICIAL INTELLIGENCE CAN REALLY DO FOR US. FOR LARGE OR SMALL BUSINESSES, AI HAS MOVED BEYOND BEING A TECHNOLOGY TO APPROACH CAUTIOUSLY. IT HAS BECOME A TOOL TO ADOPT AS SOON AS POSSIBLE TO ENHANCE EVERYDAY WORK. STUDIES REPEATEDLY SHOW THAT MANY COMPANIES ARE STILL IN THE EARLY STAGES OF AI ADOPTION. THIS IS OFTEN DUE TO A LACK OF UNDERSTANDING, UNCERTAINTY ABOUT WHERE TO APPLY IT, OR NOT HAVING THE RIGHT PARTNER TO GUIDE THEM. IN MY EXPERIENCE, DIRECTLY EXPERIMENTING WITH AN AI TOOL CAN QUICKLY DISPEL MANY OF THESE DOUBTS.

GENERATIVE AI IS PARTICULARLY SUITED TO THIS HANDS-ON APPROACH, AS IT DOESN'T REQUIRE HUGE UPFRONT INVESTMENTS OR SPECIALISED EXPERTISE. AND, AS IS OFTEN THE CASE, ONCE YOU START, YOU'LL WANT TO EXPLORE MORE.

THIS SECTION WILL EXPLORE PRACTICAL EXAMPLES OF HOW GENERATIVE AI CAN TRANSFORM BUSINESSES' CREATION, COMMUNICATION, AND OPERATION. WHETHER IT'S PRODUCING AN ARTICLE IN MINUTES, PERSONALISING CUSTOMER INTERACTIONS, OR PLANNING YOUR NEXT MARKETING CAMPAIGN, WE'LL SEE HOW THIS TECHNOLOGY CAN TRULY SUPPORT AND ELEVATE BUSINESS OPERATIONS.

Generating Textual Content.

Creating content has always been one of the most time-consuming tasks. Consider websites, blog articles, social media posts, videos, presentations, reports, or sales materials. Or think about an organisation's editorial calendar, social media strategy, and project updates. In all these areas, generative AI plays a pivotal role.

Tools like ChatGPT, Claude, or Gemini enable you to draft articles, emails, product descriptions, or social media posts in a fraction of the usual time, by providing the proper context and crafting an appropriate prompt. You can generate content that's already 80% ready, requiring only minor tweaks and validation. These tools can also analyse your existing web pages and help you reframe the text to make it more engaging and meaningful for your audience.

It's not just about speed; these tools excel at creating personalised, compelling, and audience-tailored content. They can even effortlessly translate your text into multiple languages, streamlining the process of internationalising your website or online shop when needed. This doesn't mean dumbing down your communications. You can provide precise instructions to ensure that cultural nuances are integrated into technical jargon, ensuring your message resonates effectively with its intended audience.

Generating Images.

Beyond text, some platforms enable the creation of images using natural language prompts. This means you can now complement your content with images, infographics, and illustrations tailored to your needs.

The barrier to entry has significantly lowered here as well. Instructions are provided via chat, and a vast range of prompts enables you to create images precisely as you envision them. Tools like Midjourney, DALL-E, Ideogram, and others empower users to generate visuals by adjusting parameters such as description, dimensions, filters, styles, camera settings, shot types, focal apertures, and painting techniques. Platforms like Napkin AI even allow you to select a portion of an article and generate a themed graphic to accompany that section.

This is achieved through simple prompts, which deliver immediately usable results that save considerable time while providing fresh perspectives for creative projects.

Generating Videos and Avatars.

Content generation doesn't stop at text and images. Generative AI supports various platforms for creating customised videos and virtual avatars that can directly interact with your audience. These avatars can explain product features or provide detailed assistance and can be integrated into diverse contexts, such as company websites, social media channels, or virtual events. This brings a human touch to digital interactions without requiring additional human resources.

Moreover, AI can create technical videos in one language and automatically dub them into other languages, complete with synchronised lip movements. This improves accessibility for international audiences and offers significant cost and time savings. There's no longer a need to duplicate the entire video production process for each language—professional results can be achieved with minimal resource investment.

Platforms like HeyGen, Rask, Veed, or ElevenLabs are specifically designed for these purposes and often provide free trials, making it easy to explore these capabilities without significant financial commitment. With these tools, creating video content tailored to your audience has become more accessible than ever, allowing businesses to expand their multimedia communication strategies in fresh, personalised ways. It's worth a try!

Music.

Generative AI also opens the doors to music production, offering exciting opportunities to accompany various projects. Today, many platforms allow users to create audio tracks or voices, enabling experimentation with fresh ideas while significantly reducing production costs and timelines. These tools facilitate background music creation and enable users to explore diverse musical styles, tailoring soundtracks to meet specific project needs.

For instance, you can create thematic tracks for corporate videos, soundtracks for presentations, or background music for promotional events. These musical bases can be customised in many cases, allowing users to choose instruments and rhythms and adjust tonalities to match the desired atmosphere. Achieving such levels of personalisation would have required the expertise of musicians and professional producers just a few years ago, but now, it's accessible to virtually anyone.

You can create a thematic song for your business using tools like Suno or Udio by simply inputting a brief description. You'll be amazed at the quality of these AI-generated productions. Moreover, multiple versions of the same track can be generated, allowing you to select the one that best fits your project, offering remarkable flexibility and creativity.

Strategic and Operational Support.

AI is not just about creativity; it's also a powerful tool for strategy and operations. There's a saying that you shouldn't ask AI for answers but rather perspectives. AI tools help businesses think more efficiently, quickly, or in greater detail through a data-driven or inspirational approach. AI can significantly enhance how we analyse competitors, identify target customers, generate new ideas, evaluate web pages or newsletters, develop strategies, or optimise internal processes. With access to an immense reservoir of knowledge, AI has internalised primary techniques in planning, commercial and marketing strategy, HR, sales, and operational management.

Imagine brainstorming with a "virtual colleague" that's always on. Depending on your needs, you can ask for fresh ideas to promote products or services, refine existing ones, or develop documents to support sales teams or customer service. You can even have it analyse competitor documents alongside your own to highlight differences and suggest improvements. It's possible to simulate conversations with visionaries like Elon Musk, Steve Jobs, or Jack Welch by turning them into virtual "advisors" for your business, offering input on strategies and actions.

AI can also assist in drafting briefs for agencies or consultants. By iterating with AI, you can ensure that your briefs are clear and comprehensive and include all key points your partners need to deliver their expertise effectively.

Market segmentation also becomes simpler. AI can help create detailed customer personas or Ideal Customer Profiles (ICP) with associated customer journeys. Based on these profiles, you can feed the AI a product document and ask how to position its benefits effectively. This can be done for multiple audiences, enabling you to develop tailored communication plans for each.

AI is particularly effective in data analysis. With the correct input, it can review campaign results and suggest improvements, identify unnoticed sales patterns, or uncover customer clusters with unique behaviours. While these tasks can be done manually, AI offers the advantages of speed and the ability to present different perspectives, saving valuable time and providing deeper insights into your operations.

Research.

With the integration of web search capabilities into LLMs (Large Language Models) such as SearchGPT (ChatGPT), Gemini (Google), or Perplexity, AI-powered search now complements traditional search engines like Google, Bing, or Safari.

The primary benefit is that LLMs provide comprehensive summaries of the information you seek and links for further exploration. In contrast, traditional search engines usually present a list of links, requiring the user to navigate through websites to locate the desired information.

This new research method is particularly beneficial for those who need quick, consolidated insights, such as when preparing reports or analysing market trends. AI not only saves time but also delivers a well-structured starting point, offering an excellent alternative to traditional search methods.

Conversely, you can also optimise your web pages so that AI can identify them as authoritative sources when users make related inquiries. This is a rapidly evolving area, but key practices include improving the quality of your text, crafting clear and descriptive titles (ensuring the subject is evident), and incorporating Q&A or FAQ sections into your pages. These strategies reflect the best thing to do when writing content: writing for people and providing valuable insights. Good-crafted content can make AI search engines more likely to reference your content and cite your site as a trusted source.

Voice Usage.

Integrating advanced voice systems in platforms like ChatGPT and Gemini introduces an assistant that can speak and interact with users in real-time, providing valuable support in both professional and personal contexts.

For example, these systems can be used to prepare for job interviews, presentations, pitches with investors or suppliers, and even to simulate an interview scenario. Additionally, they are excellent learning tools, enabling users to study new languages or acquire new skills with guided support that adjusts to their existing knowledge of the topic.

This ability to tailor the learning experience to the user's level of understanding makes the process more effective and personalised, ensuring better outcomes for a wide range of applications.

05

STRATEGIES FOR IMPLEMENTING AI IN BUSINESS.

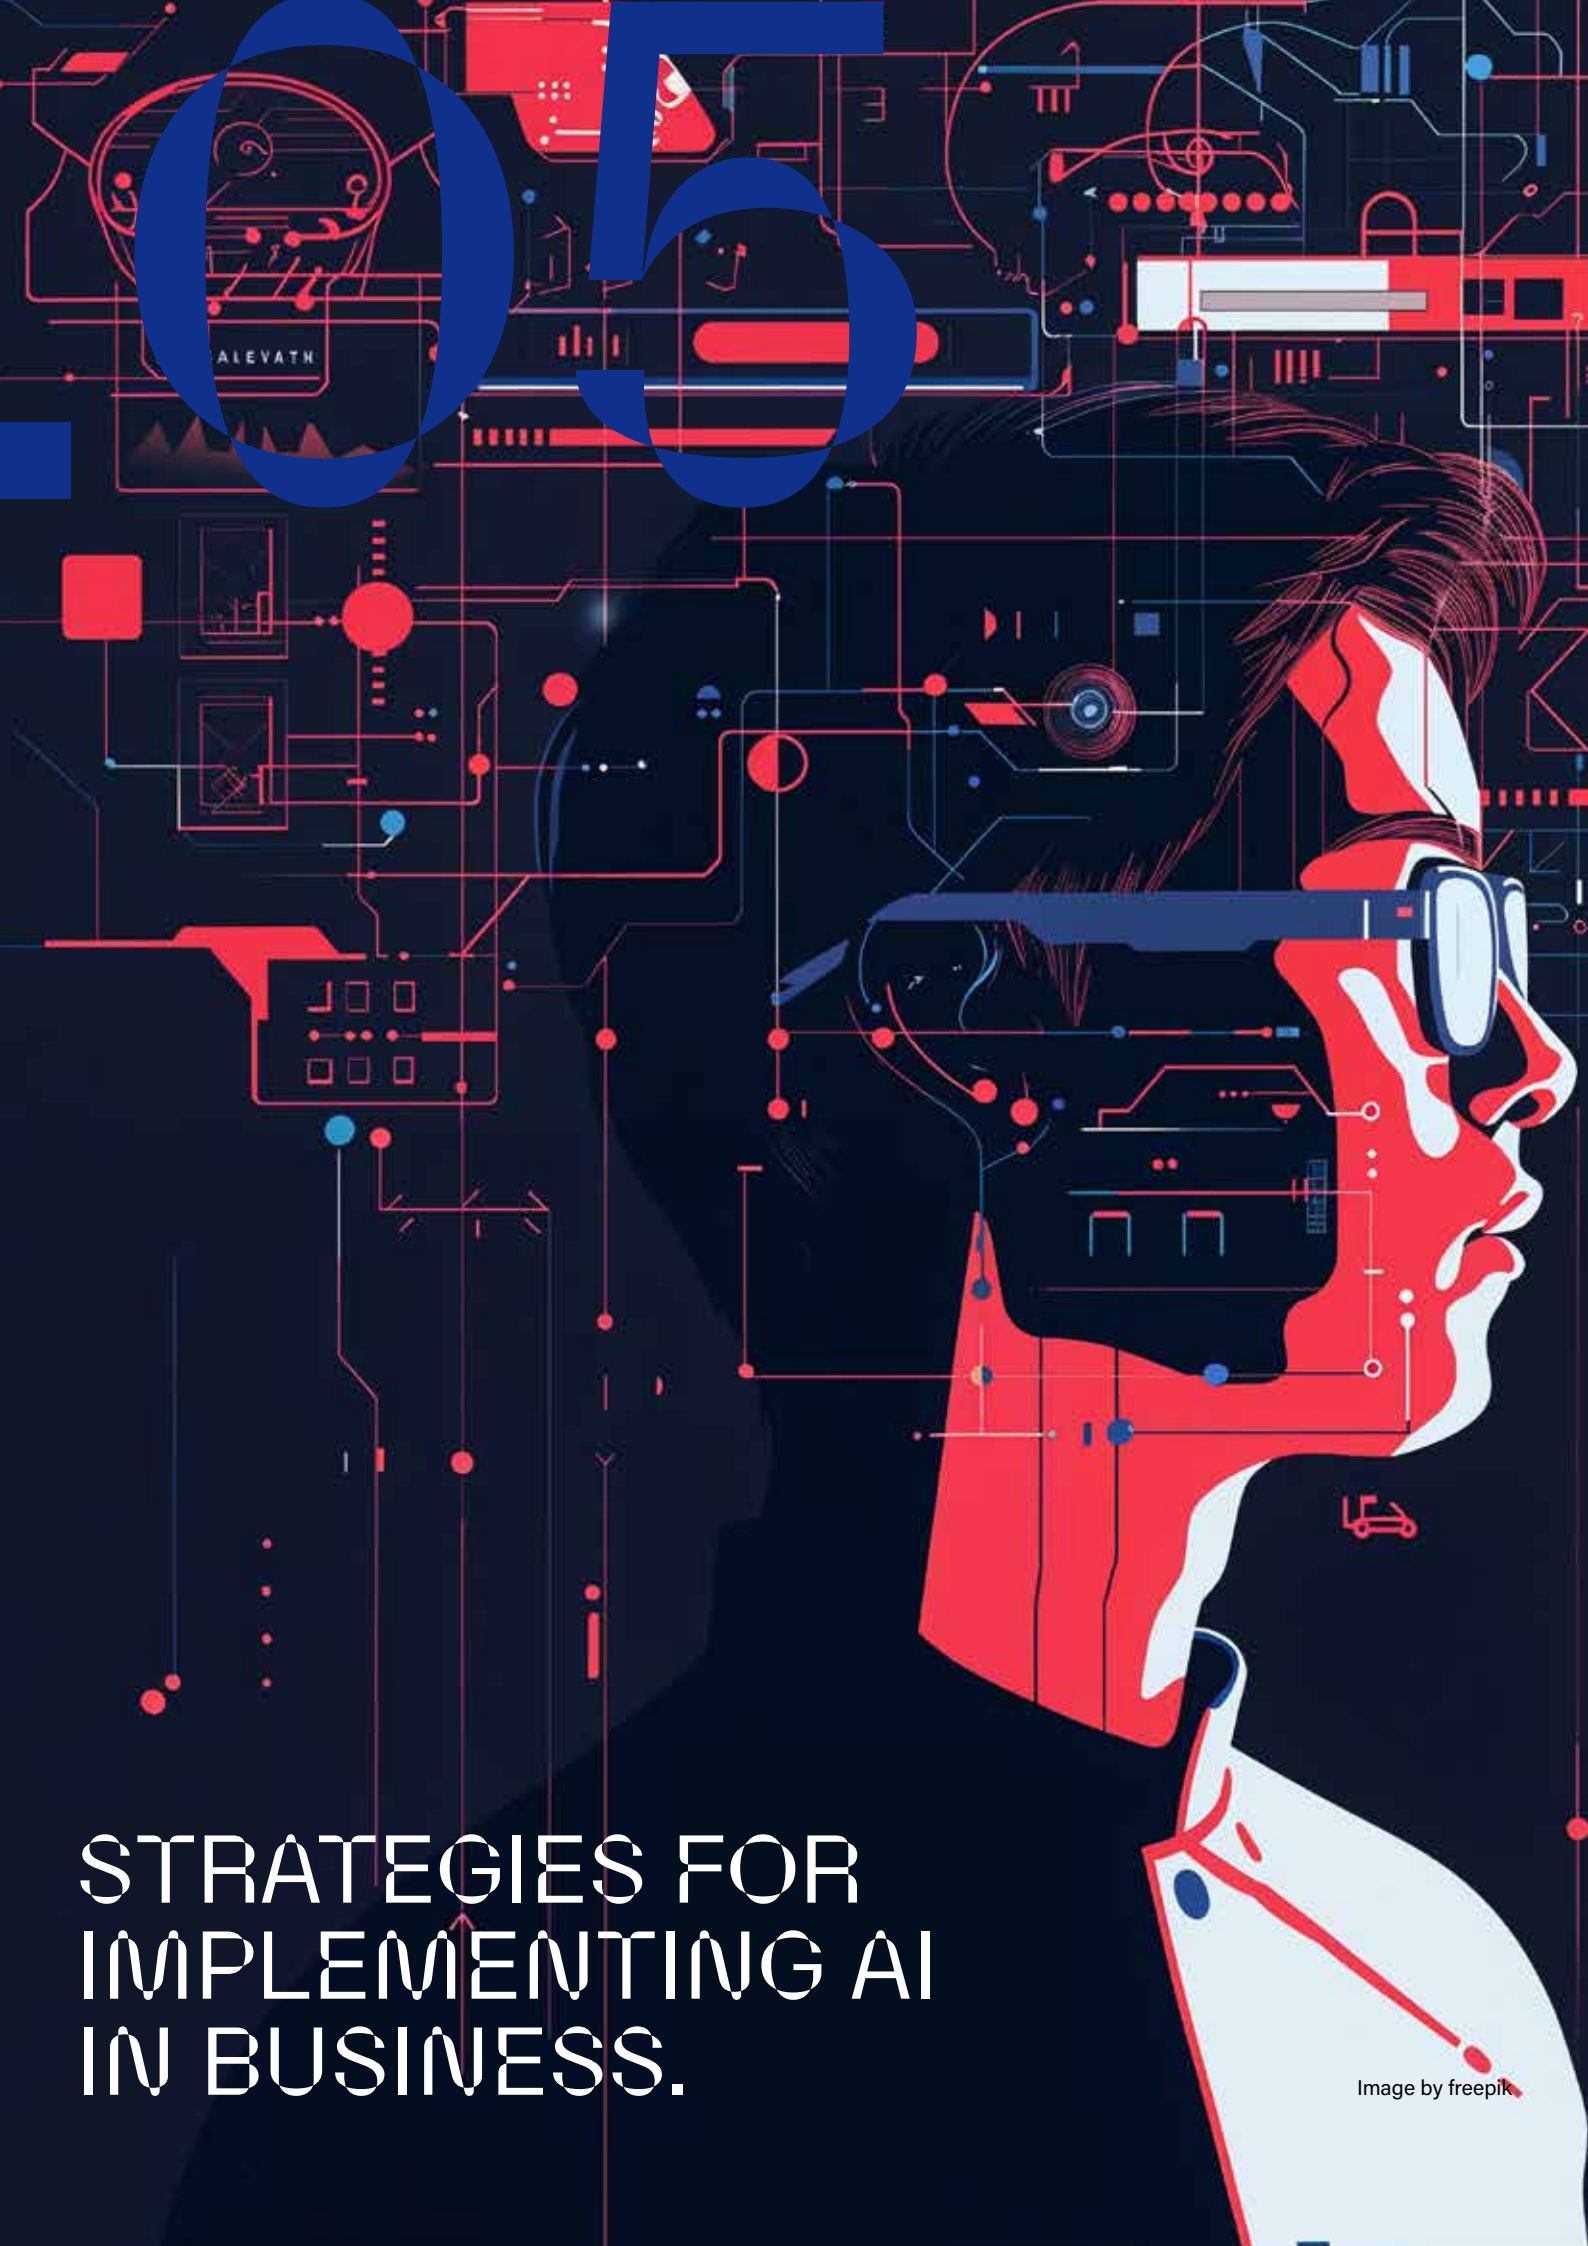

Image by freepik

ADOPTING ARTIFICIAL INTELLIGENCE ISN'T JUST ABOUT SELECTING THE RIGHT TOOLS; IT'S A TRANSFORMATION THAT REQUIRES THOUGHTFUL PLANNING AND STRATEGY. FOR THE FIRST TIME IN A WHILE, INTEGRATING THESE TECHNOLOGIES INTO KEY POINTS OF OUR PROCESSES LARGELY DEPENDS ON US.

WE KNOW OUR WORKFLOWS, WEAKNESSES, AND OPPORTUNITIES FOR IMPROVEMENT OR ENHANCEMENT. THIS INSIDER KNOWLEDGE PUTS US IN THE BEST POSITION TO IDENTIFY WHERE AND HOW AI CAN HAVE THE GREATEST IMPACT.

IT'S UNLIKELY THAT A CONSULTANT, AT LEAST INITIALLY, WOULD HAVE THE SAME INSIGHT INTO OUR OPERATIONS. THAT'S WHY THE SUCCESS OF AN AI IMPLEMENTATION STRATEGY DEPENDS MAINLY ON OUR WILLINGNESS TO TAKE AN ACTIVE ROLE AND EXPLORE ITS POTENTIAL.

LET'S LOOK AT SOME ESSENTIAL STEPS FOR CRAFTING AN EFFECTIVE ADOPTION STRATEGY.

Treat AI Like an Intern for Your Company.

I borrow this analogy from Ethan Mollick, Associate Professor at the Wharton School of the University of Pennsylvania and a recognised expert in the application of AI in the workplace, education, and daily life.

Integrating artificial intelligence into a business is like bringing a new intern into the team: eager and full of potential but needing careful guidance to turn its capabilities into tangible results. This "intern," however, has unique traits: it works at lightning speed. It has access to an astonishing amount of information. Still, it lacks practical experience, context understanding, and the ability to work independently (at least until AI evolves into autonomous agents). To fully realise AI's potential, it's crucial to understand its strengths and weaknesses, assign precise tasks, and monitor its outcomes. Like any new team member, it requires patience and time to integrate effectively into company workflows.

As with any new intern, you need to decide which tasks to delegate, such as creating standardised reports or managing automated customer responses; which to collaborate on, such as brainstorming for creative campaigns or analysing complex data; and which to keep exclusively human, such as building interpersonal relationships or making strategic decisions that require emotional intelligence and experience. Establishing an environment where AI can "grow" alongside the team is essential, using regular feedback sessions—like periodic evaluations or group discussions—to assess performance, refine processes, and improve integration.

This approach facilitates the gradual integration of AI, allowing the team to familiarise itself with the technology and gradually discover its potential. Moreover, it helps team members develop new skills, enriching the overall organisational ecosystem. Imagine it this way: every team member has their own "intern," and suddenly, your company is bustling with dozens of virtual interns. That's why it's crucial to adopt AI in a way that ensures regular feedback and a culture of collaboration.
In this light, AI doesn't replace human ingenuity; it enhances it. With a well-structured onboarding process, AI can become an invaluable ally for your company.

Identifying Areas of Greatest Impact.

The first step in implementing AI effectively is understanding where it can deliver the most value. Not all activities need automation, and not all processes will benefit equally from AI integration. Are there repetitive tasks that consume significant time? What processes could be improved in terms of speed or accuracy? Or are there areas like marketing and customer support where personalisation could make a real difference?

For example, consider automating customer responses by adopting an AI-powered chatbot. This would free up your team to focus on more complex issues. Similarly, you might use AI to generate content for your website, blog, or social media channels. Mapping out these opportunities is an essential first step.

Experimenting with Pilot Projects.

Start small with pilot projects rather than attempting to transform your entire organisation in one go. Select a single process from your mapped activities, implement an AI tool, and gather feedback to evaluate its performance.

This approach reduces risks and allows you to refine your solutions as you learn what works best for your context. For example, you might test a content-generation tool to create social media posts or use a chatbot for frequently asked questions on your website. By analysing the time saved, the added value, and other benefits, you can decide whether to extend AI usage to additional processes.

Team Training.

We have repeated that AI doesn't replace human contributions; it amplifies them. To achieve this amplification, your team needs to be well-prepared. To this end, organise training sessions on how to use these tools effectively and incorporate them into your planning. Education, from workshops on writing effective prompts to hands-on sessions with specific tools, is the foundation of effective AI adoption.

Another crucial element is overcoming resistance to change. People are naturally hesitant until they fully understand AI and how it works. Consider creating roles for internal "AI champions" who specialise in AI tools and potentially work across departments. These champions can act as a dedicated innovation lab, continuously experimenting with new applications and updating processes.

These "evangelists" will also play a vital role in communicating AI projects to the broader organisation. By gathering feedback, sharing insights, and demonstrating tangible benefits, they can help the entire company see the value of AI.

Monitoring and Optimising.

The implementation of AI is never a "finished" process. Once the tools are in place, monitoring the results and regularly identifying improvement areas is wise.
AI itself can provide valuable insights, such as performance metrics or suggestions for optimising workflows, allowing for continual enhancement.

Recognition and Rewards.

Numerous articles worldwide highlight that many individuals already use AI for personal work-related tasks. However, this doesn't automatically lead to increased productivity at the organisational level. Why not? Simply put, if such practices are not incentivised or formally recognised, expecting higher productivity levels becomes a delicate issue.
Acknowledging and rewarding the additional productivity and quality—whether through tangible rewards or symbolic recognition—can encourage widespread adoption and engagement.

Consultancy Involvement.

Experimenting with AI and running small projects is an excellent way to dispel doubts and become better equipped to engage with an expert consultant.
At some stage, seeking advice from an AI expert will be necessary to evaluate your progress, identify areas for improvement, and establish a team or framework to drive AI adoption further within your organisation.

It's important to understand that your expertise doesn't have to lie in AI itself but in using it effectively for your business goals. Therefore, a thoughtful approach to consultancy, undertaken once you've gained the foundational knowledge needed to engage meaningfully, is highly recommended.

06

ADDITIONAL ASPECTS TO CONSIDER.

Image by freepik

WHEN DISCUSSING THE ADOPTION OF AI, IT'S ESSENTIAL TO ACKNOWLEDGE SOME IMMEDIATE CHALLENGES ORGANISATIONS ARE LIKELY TO FACE. AMONG THE KEY HURDLES ARE RESISTANCE TO CHANGE, CONCERNS OVER DATA SECURITY, LACK OF TECHNICAL EXPERTISE, AND COMPLEXITIES RELATED TO INTEGRATION INTO EXISTING BUSINESS PROCESSES. LET'S EXAMINE THESE ONE BY ONE.

Fear of AI and its Complexity.

It is common to fear being replaced or learning complex new tools. To mitigate this resistance, it's crucial to communicate the goals of AI adoption, emphasising that these tools are designed to enhance, not replace, human skills. Engaging employees and management from the outset is essential to easing this transition.

Many fears will dissipate once individuals can experiment with AI themselves. Additionally, it is highly recommended that they partner with experienced consultants during this journey. Ideally, this should be done after internal experimentation, ensuring that they are more informed and can better assess the advice and proposals they receive.

Data Security.

Adopting AI inevitably raises questions about data handling and whether sharing data with AI tool providers compromises privacy. Therefore, selecting tools that adhere to the highest security standards and regulatory compliance requirements is critical.

It's also worth considering using paid versions of AI platforms, often including data encryption and confidentiality guarantees. Equally vital is educating staff on managing what files are shared with AI tools. Many AI tools use uploaded data to train their models further, potentially exposing sensitive or proprietary information.

If data confidentiality is particularly critical, companies should consider adopting proprietary platforms where data is managed and stored internally. This approach ensures tighter control and security. A skilled consultant can help you evaluate these options and select the best path for your organisation.

The Quality of Data. A Foundational Element for Business AI.

Data analysis is the driving force behind strategic decision-making, and in the context of artificial intelligence, data quality becomes an essential factor. While commercial tools like ChatGPT can generate value without direct access to company-specific data (you could, but it is out of scope in this guide), the situation changes significantly when integrating a proprietary AI system. In such instances, the AI must interact with sources such as CRM, CDP (Customer Data Platform), analytics data, ERP, and other enterprise platforms.

Data quality is non-negotiable in this context. Incomplete or inaccurate data can lead to flawed analyses, incorrect predictions, and poor strategic decisions. For instance, if a CRM contains duplicate records or outdated information, the AI might personalise marketing campaigns for clients who are no longer relevant or even exist.

Ensuring data quality requires implementing practices such as verifying data sources, regularly cleaning databases, and normalising information from various systems. These actions provide a solid foundation for AI systems and enable the creation of an integrated ecosystem that ensures consistency across enterprise platforms. This is where proprietary AI can shine, offering accurate predictive analyses, logistical optimisations, and tailored insights based on reliable, up-to-date data.

In summary, while starting with commercial tools is an excellent way to explore AI capabilities, focusing on data quality is critical for developing customised systems that meet specific business needs and deliver tangible value in the long term.

Publishing Validated Content.

While AI's ease of creating content is a significant advantage, validating such outputs to ensure they align with the company's tone, values, and factual accuracy is essential. Publishing content with errors, inaccuracies, or statements that violate social norms can significantly damage a company's reputation.

Generative AI systems are trained on extensive datasets sourced from the web and publicly available databases. Human reviewers frequently refine and optimise their outputs through techniques like Reinforcement Learning. However, if the training data contains errors, biases, or inaccuracies, these issues will inevitably surface in the AI-generated content.

Thus, expert validation is crucial to maintaining credibility and safeguarding the company's image. It ensures that the final output aligns with corporate values and meets professional standards.

Humanised AI, but Not Too Much.

As advanced as AI may seem, it doesn't "think" like a human or possess sentience. However, we often tend to attribute human characteristics to it (a process known as anthropomorphising). Its responses are based on statistical correlations between words and concepts, not on true understanding. This is a crucial point to always bear in mind. Treating an AI model as a human interlocutor can improve interactions. It can make dialogues smoother and even help refine your thoughts when crafting your requests.

It's important to emphasise that current AI models do not possess abstraction or deduction capabilities like humans. However, researchers are observing surprising phenomena referred to as emergent capabilities. These unplanned behaviours appear in AI models as their complexity and the number of parameters increase. While not intentionally programmed, these abilities border on abstraction and deduction — for example, solving logical problems, translating intricate texts, or adapting to new tasks with minimal examples.

These behaviours likely arise from the non-linear interactions among the billions (or even trillions) of parameters in AI architectures. However, the internal workings of these systems are so intricate that even their developers often cannot fully comprehend them. Emergent capabilities represent both a challenge and a marvel of modern artificial intelligence: an uncharted frontier where each new observation raises fresh questions and opportunities for the future.

THAT'S A WRAP!

Today's artificial intelligence tools are powerful, accessible, and require little investment. However, no AI can replace the context, strategy, and vision of a company's leadership. Implementing AI requires effort and responsibility, but the results can remarkably transform a company's future.

There's a phrase often repeated: "AI won't take your job, but someone using AI likely will." To rephrase: "The competition won't take your market, but a competitor leveraging AI probably will." This isn't a call to fear but a prompt to take action.

Instead of artificial intelligence, consider augmented intelligence. Generative AI tools don't replace who we are but amplify our capabilities, enhancing our creativity and problem-solving skills.

Think of it as an additional layer for the mind that magnifies that human spark within us — our expertise, sensitivity, awareness, and practicality — enabling us to create and achieve things we never imagined. We won't lose our humanity; instead, we'll have more time to think critically, pursue exciting projects, and connect more meaningfully with others.

Now is the perfect time to begin. Once you experiment with AI, you'll discover a new way of working that you won't want to abandon. Here's to augmented intelligence!

Disclaimer

The views expressed in this guide reflect the author's personal perspectives and are intended solely for informational and educational purposes. The author has no commercial or promotional intent beyond providing insights into the evolving landscape of artificial intelligence. Any trademarks or brand names mentioned belong to their respective owners, and their inclusion in this guide does not imply endorsement or affiliation.

Glossary

This brief glossary is designed to help you familiarise yourself with key artificial intelligence (AI) terms. While you don't need to memorise them immediately, having them on hand will help you navigate the AI landscape more confidently.

- **Algorithm**
A sequence of instructions a computer follows to perform a specific task. In AI, algorithms are the foundation for learning and generating results.
- **Agent**
An autonomous software entity that uses AI to carry out specific tasks based on assigned objectives. Agents can interact with their environment, make decisions, and adapt to changes. Examples include chatbots that answer customer queries or recommendation systems suggesting products.
- **Supervised Learning**
A type of machine learning where the model is trained using labelled data, meaning the correct answers are already known. For example, images of cats labelled as "cats" help the model recognise them in future.
- **Unsupervised Learning**
Unlike supervised learning, the model works independently with unlabelled data, discovering patterns and correlations. It can help cluster similar data or uncover hidden relationships.
- **Bias**
Prejudices or distortions in the data used to train AI can lead to inaccurate or discriminatory outcomes. For example, a hiring system that inadvertently favours one gender over another.
- **Chatbot**
A virtual assistant using AI to answer questions, provide information, or assist customers. Examples include ChatGPT and customer support bots on websites.
- **Deep Learning**
Deep learning is an advanced subset of machine learning based on deep neural networks. These models can process large volumes of complex data, such as images, videos, and language, and deliver highly accurate results.
- **Diffusion Model**
A type of model used in generative AI, particularly for creating images. It starts with random noise and iteratively "reconstructs" an image, gradually improving its quality. This approach powers tools like DALL-E and Stable Diffusion.
- **Eliza**
One of the first chatbots, based on conditional rules, was created in the 1960s. It simulated a conversation with a psychotherapist by responding with generic or pre-set phrases—a classic example of early, limited AI.
- **Embeddings**
Numerical representations that transform words, phrases, or other data into vectors, enabling AI models to analyse and compare meanings. They're used to grasp the semantic context in natural language or identify similarities between concepts.
- **Generative AI**
AI that goes beyond data analysis to create original content such as text, images, music, or video. Examples: DALL-E (images) and ChatGPT (text).
- **Conditional AI**
A more straightforward form of AI based on predefined rules, where specific inputs determine actions. For instance, "If the user writes X, respond Y". These systems paved the way for more advanced technologies.
- **Predictive AI**
AI that analyses historical data to forecast future events. It's commonly used to predict product demand, identify machinery faults, or anticipate financial fraud.
- **Large Language Model (LLM)**
Generative AI models trained on vast amounts of text to understand and generate natural language. Examples include ChatGPT and Bard.
- **Machine Learning (ML)**
A subset of AI where computers learn from data without explicit programming. It allows systems to improve their performance by "learning" from experience.
- **Overfitting**
A machine learning error where the model "learns too well" from training data, making it unable to generalise effectively for new data.
- **Prompt**
Instruction or requests are given to an AI system to generate a response or result. Clear, detailed prompts are crucial for quality outputs.
- **Neural Networks**
Models are inspired by the human brain's structure, which consists of layers of nodes (artificial neurons) that analyse data to recognise complex patterns.
- **Tokens**
AI models process data using the smallest units of language. Depending on the context, a token can be a whole word, part of a word, or a symbol. For instance, "Cat" is a token, but "Ca—" and "t" could also be tokens in some models.
- **Training**
The process through which an AI model "learns" from provided data, adjusting its parameters to produce increasingly accurate results.

A SHORT GUIDE TO

ARTIFICIAL INTELLIGENCE

FOR BUSINESSES AND ENTREPRENEURS

VITTORIO NERI

Con oltre 20 anni di esperienza nel marketing B2B, oggi mi occupo di pianificazione, ricerca e tecnologie di marketing per Roland DG. Il mio obiettivo è migliorare le strategie aziendali e valorizzare prodotti e servizi che facciano davvero la differenza per le persone. Ho avuto la fortuna di crescere professionalmente tra Italia, Europa e Giappone, e questo mi ha permesso di specializzarmi in strategie di marketing su misura per aziende e utenti. Amo lavorare sul campo e con le persone. Le relazioni, le storie e le emozioni dei brand mi appassionano, ma sempre con un occhio attento ai dati e alla concretezza. Sono anche molto interessato alle tecnologie di marketing e alle piattaforme che migliorano il lavoro quotidiano, e ho esperienza nel collaborare con team internazionali e culture diverse.
Mi tengo sempre aggiornato, ma fuori dal lavoro non rinuncio alle mie passioni: musica, cinema, giardinaggio, viaggi e lettura.

Progetto grafico cortesemente realizzato da: Hammer Design Brand Experience