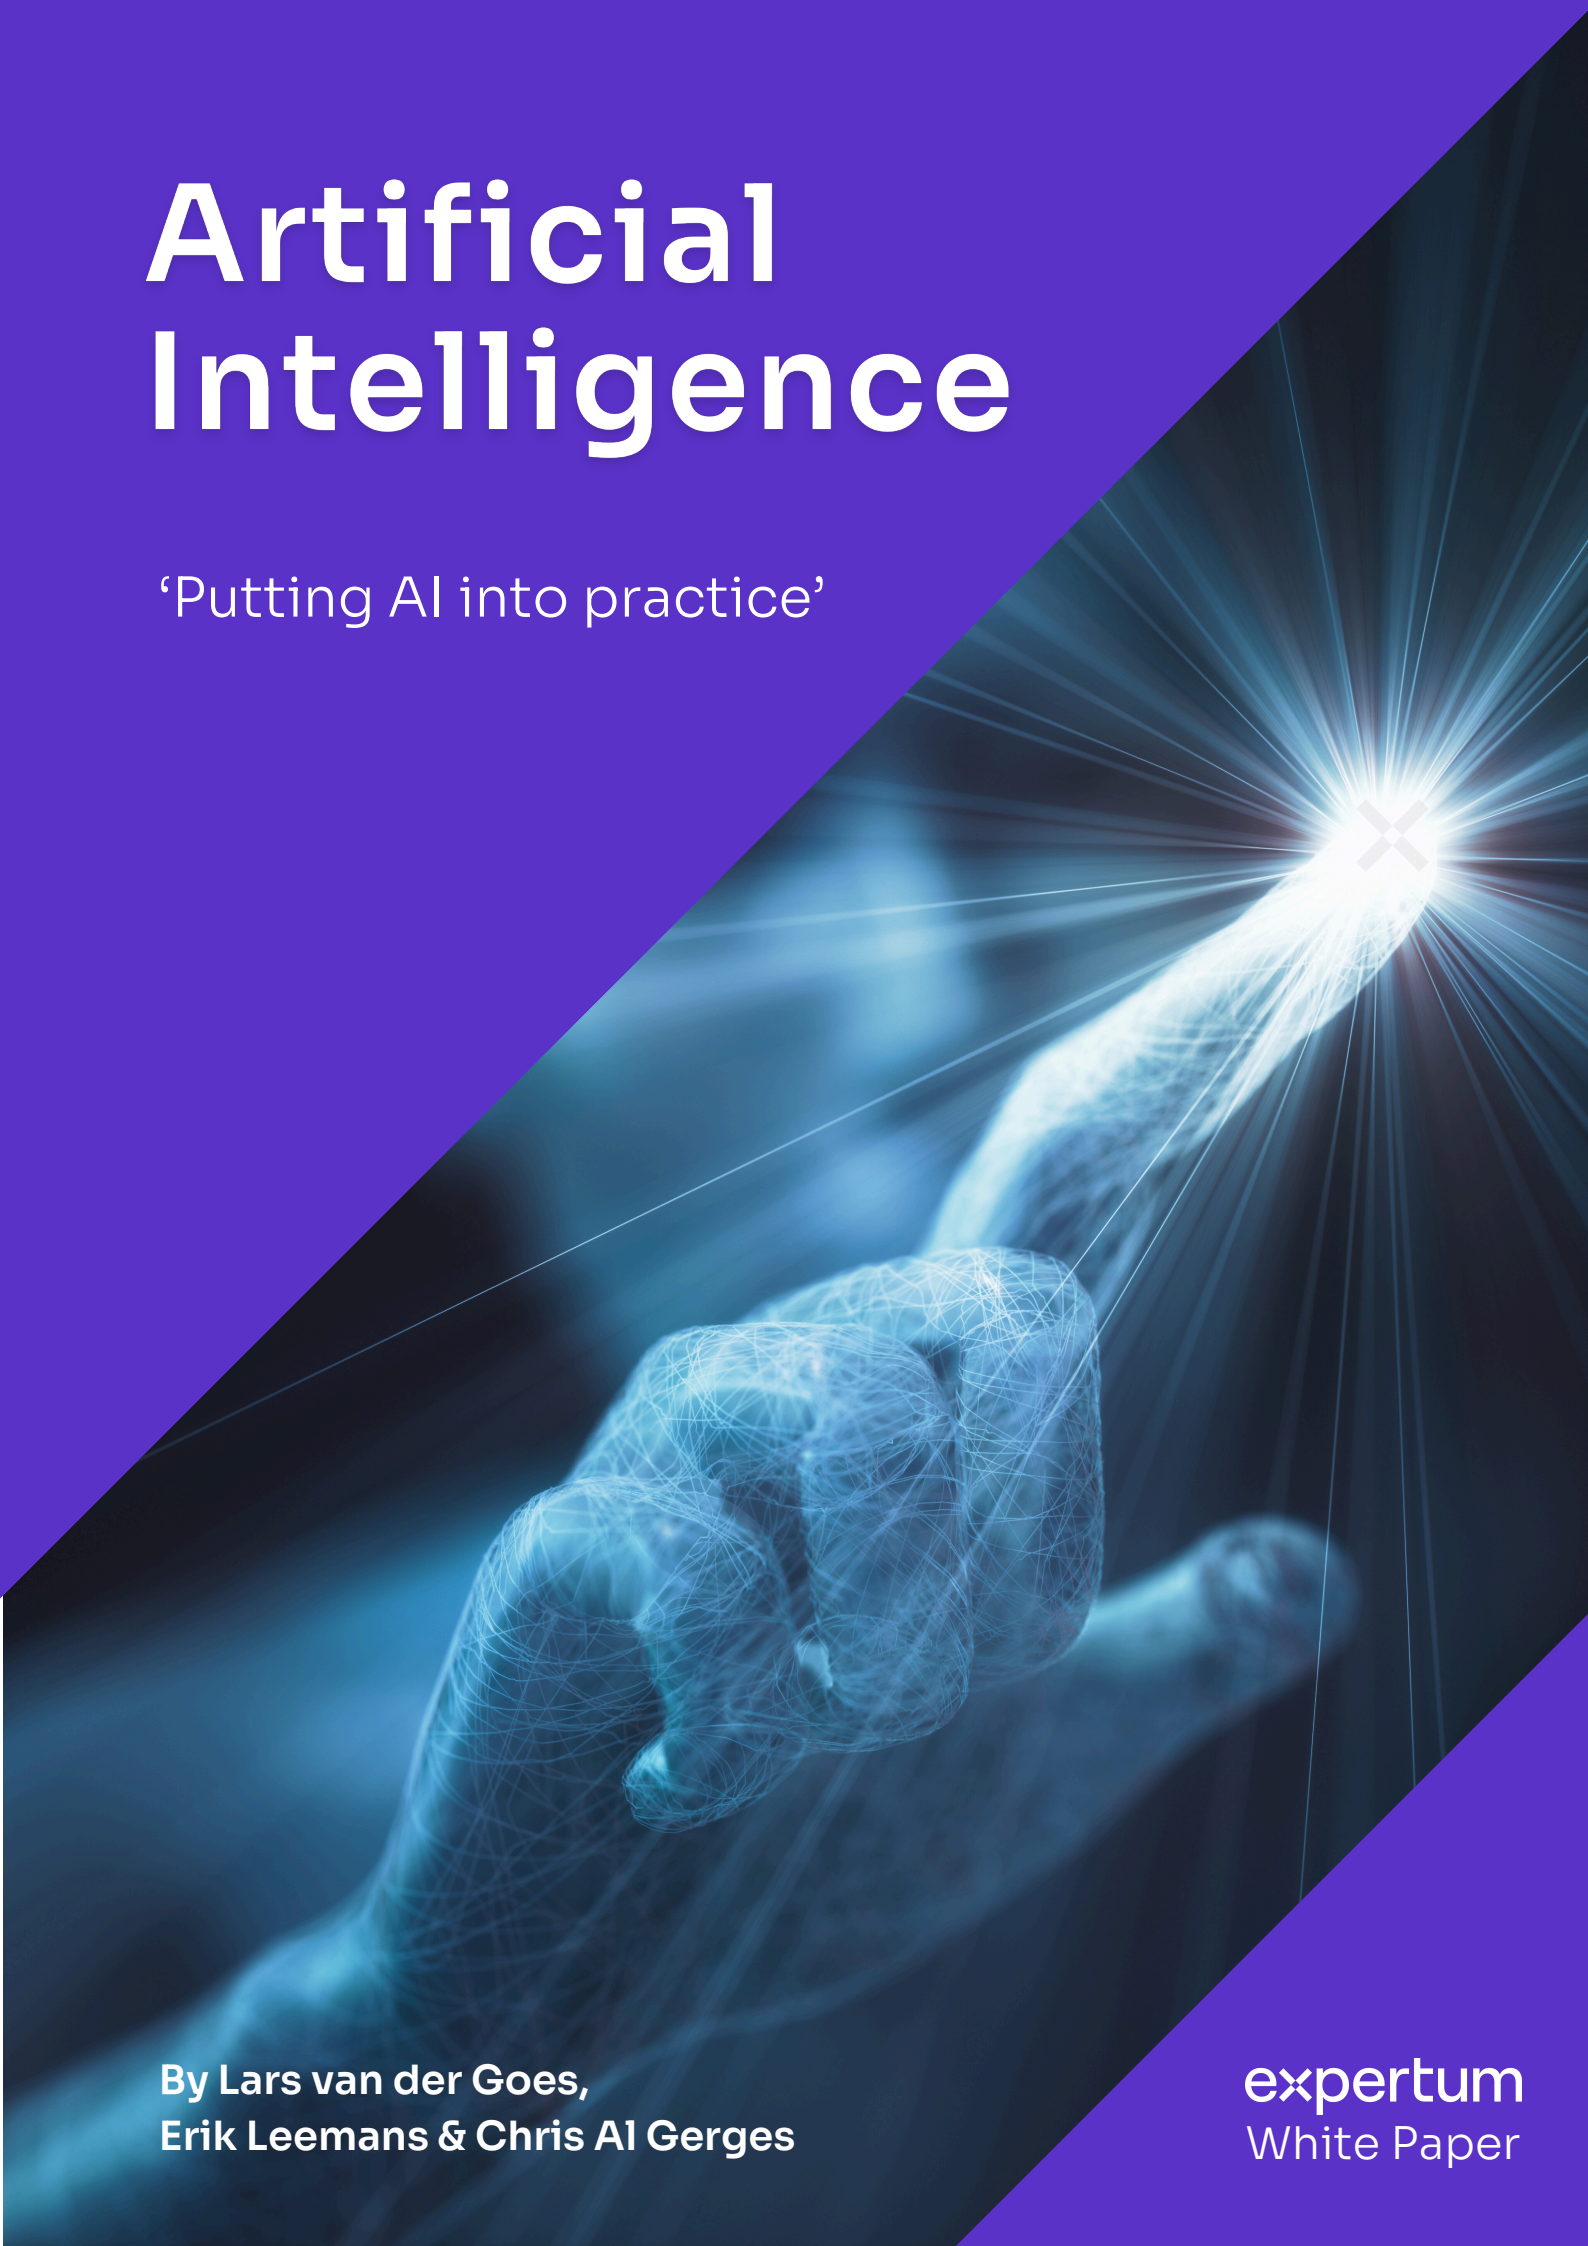
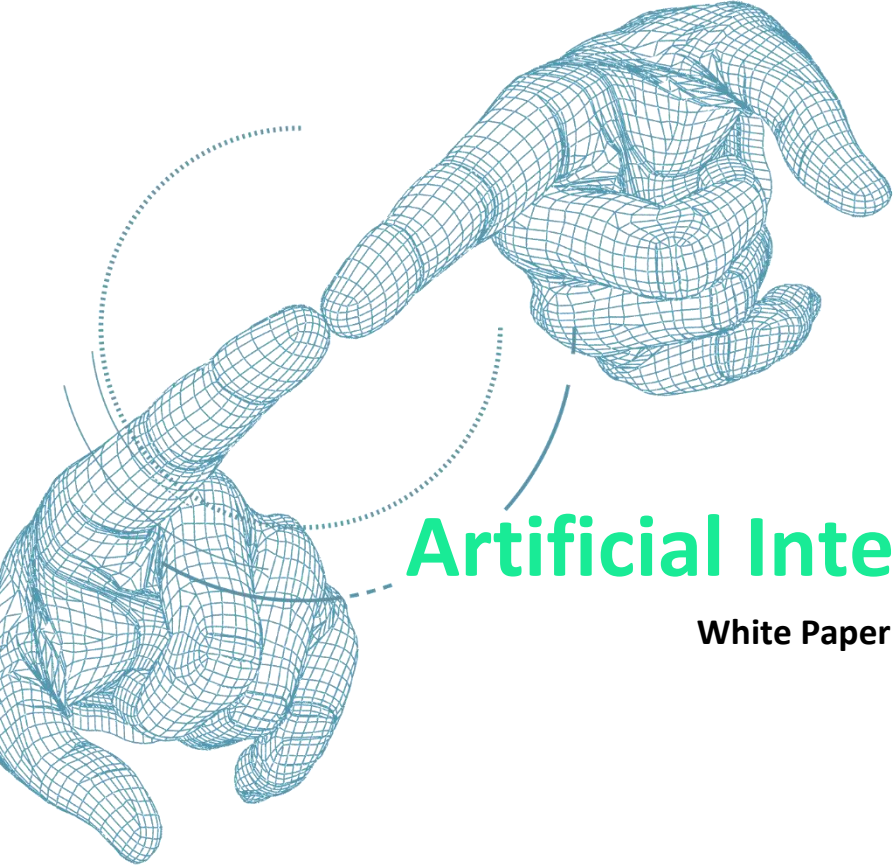
Artificial Intelligence

‘Putting AI into practice’



By Lars van der Goes,
Erik Leemans & Chris AI Gerges

expertum
White Paper



Artificial Intelligence

White Paper

V 1.0 January 2025

Lars van der Goes, Erik Leemans & Chris Al Gerges

Table of contents

Preface 3
What is Artificial Intelligence? 4
A brief history of AI 4
Understanding the fundamentals 5
The market 6
Artificial Intelligence & SAP 7
Putting Artificial Intelligence into practice: our vision 9
A practical example 10
Closing remarks 11

Preface

With an origin dating back all the way to the 1940's and involving names like Alan Turing and John McCarthy, Artificial Intelligence (AI) itself is far from new to the world. It was, however, new to the masses and our society at large when OpenAI launched its GPT-3 Large Language Model in 2020 amidst the global COVID-19 pandemic. Ever since this 'AI boom', it has become abundantly clear that Artificial Intelligence is here to stay.

For both recreational and professional purposes, AI has become impossible to omit when we think about our future as human society. The streamlining, optimizing and automation of all sorts of processes is taking exponential leaps, as the way our world works changes at an ever-increasing pace. Estimates have the financial footprint of AI initiatives around *15 trillion* dollars by the end of 2030[1]; a staggering amount of money and global impact.

However, with all the possibilities AI offers, it also comes with a slew of challenges and new dilemma's that we have to face[2]. From ethical considerations to privacy concerns and from adoption resistance to security risks; the (global) community has to come up with solutions and coping mechanisms to ensure that AI is kept in check and the technology is leveraged through responsible means[3]. This is especially important for organisations, be it commercial or non-profit, as they form the backbone of our society. Employing AI both effectively and responsibly thus requires a certain degree of understanding of the underlying principles of AI technology, as well as of the purposes for which it might be utilized.

In this white paper, we will have a look at the general state of AI today and zoom in on its potential within the organisational IT landscape. We will briefly elaborate on AI's history, fundamental rationale, and why we believe it is important to have at least a basic understanding of these principles. As SAP specialists, we will then dive into SAP's ambitions with the technology and explore what AI can do for (SAP) IT, before briefly touching upon some practical scenarios where Expertum can help realize AI-driven solutions *today*. Finally, we will close with our vision on and recommendations for employing AI within organisations in an effective, sustainable and responsible way.

[1] Source: PWC. (2024) 'Exploiting the AI Revolution: Analysis'.
[2] Source: Forbes. (2024) 'AI Predictions for 2030'.
[3] Source: T-G. Nicolâescu (2024). 'How Will AI Affect Jobs: How Many Jobs Will AI Replace by 2030?' *Pandectele Romane 115.*

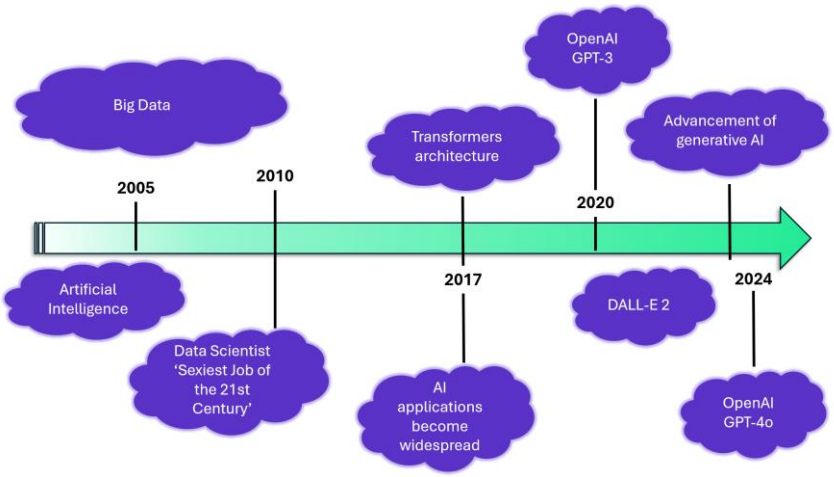
What is Artificial Intelligence?

A brief history of AI

Artificial Intelligence as we know it today finds its roots in the 1940s, when scientists from various disciplines started to think about the concept of '**machine intelligence'**. Alan Turing, known for the infamous Turing Test, was amongst these pioneers. The 40's were also the years where the first theories on artificial neural networks were penned, most notably by Walter Pitts and Warren McCulloch. What followed was a period of 'small' steps involving cybernetic robots, game AI and the cognitive revolution, which also had far-reaching implications for the fields of computer-and-data science.

After a period of great strides involving leading scientists such as John McCarthy and Marvin Minsky, among which was the first single-layer **neural network** (the perceptron), the first 'AI Winter' hit around 1974. Various challenges such as financing, the lack of sufficiently large data volumes and the limits of computing power, caused the field of AI to become dormant. The global adoption of 'Expert Systems', programs able to answer specific domain questions through logically set-up rules, reinvigorated AI in the 1980s. The fruition of probabilistic reasoning and reinforced learning soon followed, after which the second 'AI Winter' took hold towards the end of the 80s (for similar economic and technical causes as the first winter).

Finally, as computing power started to increase exponentially in the light of Moore's Law (both CPUs and GPUs), AI took precedence once again with the emergence of Big Data in the 2000s. The concept of **Deep Learning** was developed around 2012, and research into AI spiked once more. Subsequently, the invention of the transformer architecture in 2017 and the development of Large Language Models (LLMs) such as GPT-3 in 2020, signalled a new era in the age of not just AI, but human technology in general. AI became more visible and accessible to the mainstream through tools like ChatGPT, DALL-E and Midjourney. Today, the pitfalls (e.g. hallucinations[4]) and limitations (e.g. inference speed[5]) of transformer-based models are rapidly being tackled by a growing community of AI engineers across the globe.



The emergence of AI in the 21st century.

[4] https://neptune.ai/blog/llm-hallucinations
[5] https://neptune.ai/blog/transformers-key-value-caching

Understanding the fundamentals

At Expertum, we believe in understanding the workings of the solutions we employ. While this is of course a tall order when it comes to something as complex as Artificial Intelligence, we do think it is useful to at least have a grasp of the fundamentals. After all, would you rather have a partner that deploys AI solutions in your landscape and describes it as a black box that 'just works', or a partner that at least understands the logic and code that drive these solutions?

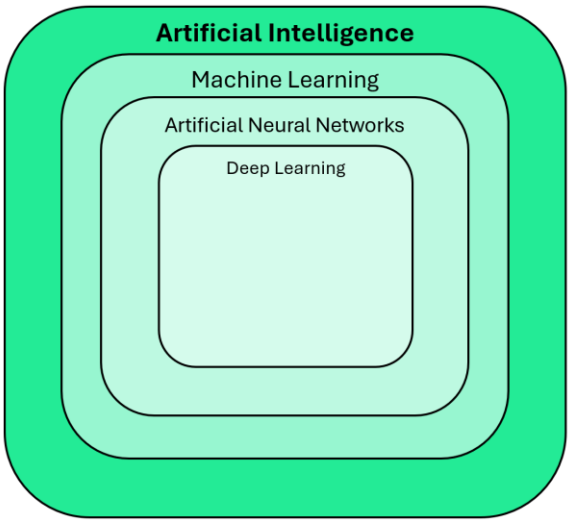
At its core, Artificial Intelligence utilizes a combination of two things: large amounts of **data** and logic in the form of **algorithms** to produce a certain output.

The first 'ingredient' is quite easy to define: this simply entails huge amounts of data, both structured and unstructured, coming from any source. Today, this can range from anything like organisational Enterprise Resource Planning (ERP) systems and databases to scanned images or social media posts. In the majority of cases, the most important thing here is that data is properly cleansed and structured before it can be used in conjunction with AI. Undoubtedly, you will have heard of the term **'garbage in, garbage out'**; this applies in particular to data in AI context. A sound data and analytics infrastructure is therefore the foundation of any successful AI implementation.

The other 'ingredient' required to realize Artificial Intelligence is logic, more specifically an algorithm. Algorithms come in many shapes and forms; you might have heard of time series analysis (with or without exponential smoothing), analysis of variance (ANOVA), linear regression (used to predict continuous results), logistical regression (used to predict binary results, or classification) or even LASSO (prevents overfitting in regression). However, modern AI applications build on the workings of these algorithms and go one step further by using **Machine Learning**, a subdomain of the AI field.

An example of a Machine Learning algorithm is a neural network. This neural network works similar to how human brain works (hence its name). A more advanced (read: deeper layered) version of a neural network is Deep Learning, which is employed by aforementioned Large Language Models such as ChatGPT. To explain the workings of a neural network briefly, we will walk through a simplified AI model development process below. If you are interested in the development process of an AI (or more generally, a data science model), look up the CRISP-DM methodology for more details.

First, the neural network (AI model) combines several (weighted) decision trees and stacks them on top of each other, creating a massive decision tree. The AI model then runs through this tree numerous times and 'measures' its own performance in terms of how well its prediction was. It uses a subset of data, called the training set, to move through this process. It subsequently validates these outcomes against a previously 'checked-to-be-correct' validation subset of the data. This way, the AI 'learns' what decisions lead to what outcomes. These are what are called the 'training' and 'testing' phases, which prepare the AI model for its intended purpose. Afterwards, the AI model can be deployed for productive use. Note that once an AI model is deployed, it can no longer 'learn'; for a more up-to-date version, you will need to build and train a new (version of a) model.



Infographic: the (sub)domains of AI shown as layers.

The market

Now that we have covered the basics of what Artificial Intelligence is, we can review the current state of the market. Today, there is a plethora of **Large Language Models** available for both commercial and private use. The most well-known include GPT (OpenAI), Gemini (Google), Llama (Meta), Claude (Anthropic) and Phi-3 (Microsoft). The access to these models also differs per case, with some being open (source) and others only being accessible via API or chatbot.

If we then shift our focus to the corporate market, we see various parties implementing AI in different ways. Cloud computing solutions offer various ways to train, test and deploy (pre-trained) AI models for specific purposes. Amazon's market-leading AWS, Microsoft's Azure and Alphabet's Google Cloud are all examples of places where some of the aforementioned models can be (partially) utilized. In practice, these applications often translate to process-supporting tools tailored to the average user or consumer. Examples include virtual assistants, chatbots and workflow automation.

Moving towards proprietary solutions, we see names such as Databricks, Snowflake, BigQuery, Redshift and Cloudera. All of these solutions offer both data management (warehousing) tools, which are crucial for proper data preparation, as well as a variety of AI functionalities to their customers. Moreover, there are data intelligence and governance tools such as Collibra, Azure Data Catalog and Atlan that will all play an increasingly important role in not just data, but AI-driven processes in the years to come. This goes double if we consider the constantly changing (and lagging!) **legislation** with regards to AI across different parts of the world, with the EU AI act[6] being a prime example of such regulation.

Finally, as the biggest ERP software vendor in the world, **SAP** could of course not shy away from implementing the potential of AI in its portfolio. Given that SAP is also where most of our expertise lies, let us have a more detailed look at how the experts from Walldorf integrate Artificial Intelligence in their solutions.

[6] https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence

Artificial Intelligence & SAP

Like many big tech companies, SAP started throttling its AI-focused developments after GPT3 was launched in 2020. The roadmaps for many of its flagship solutions, including SAP S/4HANA, SAP BTP and SAP Datasphere, even shifted to accommodate the big swing to AI-driven solutions. It became clear that AI would be a mainstay across SAP's solution portfolio going forward. We will briefly discuss the most important of these AI-driven solutions in this section.

Overall, SAP pushes the message of a portfolio-spanning **Business AI**, coupled with the values of Relevancy, Reliability and Responsibility (the three Rs). Concretely, this entails various 'small' AI applications across the various business fields in which SAP's applications are being used and embedded out-of-the-box. From finance to logistics and from marketing to customer service; AI ties these processes together. Among others, examples include error detection and resolution, automation of tasks/transactions, pattern prediction and recognition, decision planning, text analysis and of course, an embedded smart assistant called **Joule**. As SAP's AI copilot, Joule is Walldorf's primary 'mascot' for its AI efforts. Built on collaborative multi-agent systems, Joule is being integrated into virtually every solution and receives frequent updates for new functionality, extensibility options (e.g. Joule Studio) and increased language support[7].

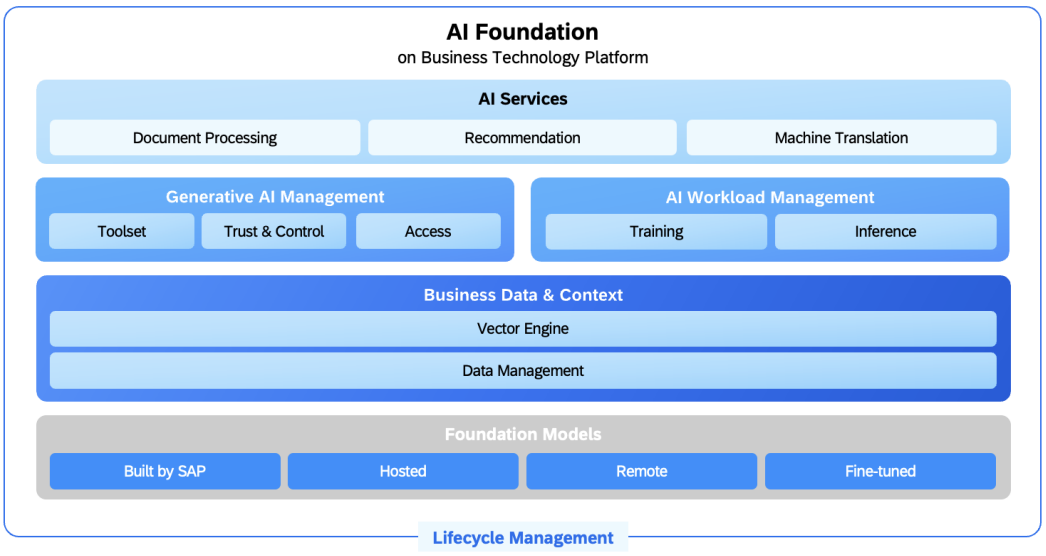
While perfectly functional in most cases, the lion's share of these AI applications is also relatively conservative. For the past few years, end users have already been using functionalities powered by AI algorithms in much of their daily lives; it was just that the 'AI-label' was not prevalently present until more recently. However, there are also areas where SAP is employing a more rigorous approach towards AI. One of these areas is the SAP Business Technology Platform (BTP).

The **AI Foundation** is SAP's all-in-one AI toolkit that is available on the SAP BTP. This solution is more tailored to AI developers, providing built-in AI services and tools for customization across the entire AI lifecycle, allowing developers to integrate AI components in business applications. Additionally, SAP allows its customers to leverage licensed versions of leading LLMs (such as GPT) by offering SAP-trained **Foundation Models** through the BTP for customer developers to then employ as they see fit. Major solutions that can be found within the BTP AI Foundation include:

- **SAP AI Core**: a service on the BTP that handles the execution and operations of an enterprise's AI assets (i.e. datasets and model files). Specifically, developers can execute and deploy AI workloads in SAP AI Core, such as model training, model inference and data processing jobs. The AI Core also comes with software development kits (SDKs), through which AI workloads can be customized.
- **SAP AI Launchpad:** a software-as-a-service (SaaS) application on SAP BTP that allows users to manage their AI scenarios across multiple AI runtimes (e.g. the SAP AI Core). Essentially, the SAP AI Launchpad provides users with graphically-based access to SAP AI Core.

[7] https://www.sap.com/assetdetail/2024/06/36693df6-c17e-0010-bca6-c68f7e60039b.html

- **HANA Vector Engine:** a functionality available in SAP HANA Cloud that allows you to save embeddings (numerical representations of data, such as text or images produced by LLMs). It also comes with special functions to process embeddings, such as similarity functions that can be used to perform a similarity search across the embeddings. The Vector Engine is also the main component responsible for Retrieval-Augmented Generation (RAG) for LLMs used within the SAP ecosystem[8].
- **Generative AI Hub:** provides generative AI models under foundation-models managed by SAP AI Core. These are executable as serving templates; developers can also use these models to enrich their applications. Examples include IBM, Bedrock or Vertex.
- **SAP Knowledge Graph:** a graphical network that provides semantic context on SAP's structured data objects and an organisation's business data, including ABAP tables, CDS views and data fields. The Knowledge Graph essentially functions as a bridge between structured data and natural (human) language.



An overview of SAP's AI Foundation.

In general, we believe that key differences can be made through SAP's foundational approach to AI, rather than its embedded applications. While the latter offer seamless experiences with AI-supported processes in operational environments, tailored, flexible and business-specific solutions can be generated by leveraging the AI Foundation.

Of course, a large part of this advantage does not come out of the box and will have to be developed by the customer natively. The AI Foundation provides the tools and (sometimes) the recipe; it is **up to the customer to decide** how to leverage these assets in the context of his/her own organisational data and processes. Although this is often not an easy process, if planned and monitored well, it can definitely provide a valuable step towards new efficiencies and competitive advantage. Expertum has done its homework, and we have come up with a clear vision, approach and practical results backed by thorough understanding to help our customers realize AI solutions the right way.

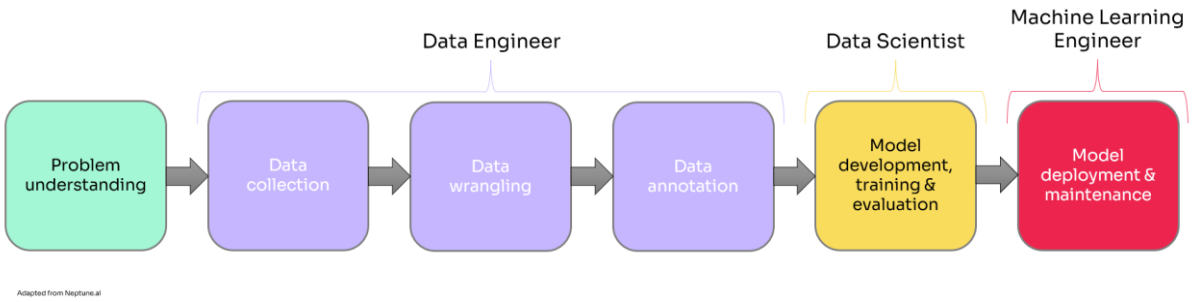
[8] https://help.sap.com/docs/hana-cloud-database/sap-hana-cloud-sap-hana-database-vector-engine-guide/introduction

Putting Artificial Intelligence into practice: our vision

Despite the large influence of statistics in AI, the more complex the model, the more difficult it is for humans to deduce exactly how the model came to certain conclusions. This lack of transparency combined with the currently still low level of in-depth domain knowledge on data science, makes it difficult (and in some cases even risky) for organisations to integrate AI solutions in their key processes. An additional challenge in AI projects is the breadth and depth of such a project; it is often easy to get drawn into pure mathematical model evaluation and lose sight of the **business case** at hand. This in turn leads to over-scoping and overspending, which are words no decision maker ever wants to be faced with.

So how can we mitigate these pitfalls and work towards efficient *and* effective AI solutions? To start, every AI project requires **specific expertise** on top of the relevant business knowledge for the challenge at hand. Data scientists have long been the most sought-after experts in the global market, but ever since Large Language Models and commercial Foundation Models (pre-trained AI models that can be tweaked for specific goals) have shown up, the need for Machine Learning Engineers has been the new 'trend'. The reason for this is that anno 2025, most large companies have their own in-house data scientists. Hence, conceptually realizing a fitting model is no longer the biggest challenge; the challenge now is to effectively deploy and maintain said model.

At Expertum, we possess extensive technical and functional solution knowledge, as well as experience with SAP systems across the portfolio. Additionally, we have the foundational knowledge in data science and mathematics to ensure that AI solutions are also vetted in this regard. We do this by clearly defining functional (business) targets, having a proper overview of and grip on relevant data and by sticking to a tailored project-and-governance plan. We also believe in the clear but **cooperative role separation** that is subscribed to by many experts in the field[9], as we employ data engineer, data science and platform expertise to train, test and deploy an effective AI model within an SAP landscape. With our vision and approach clear, let us discuss a real-life use case.

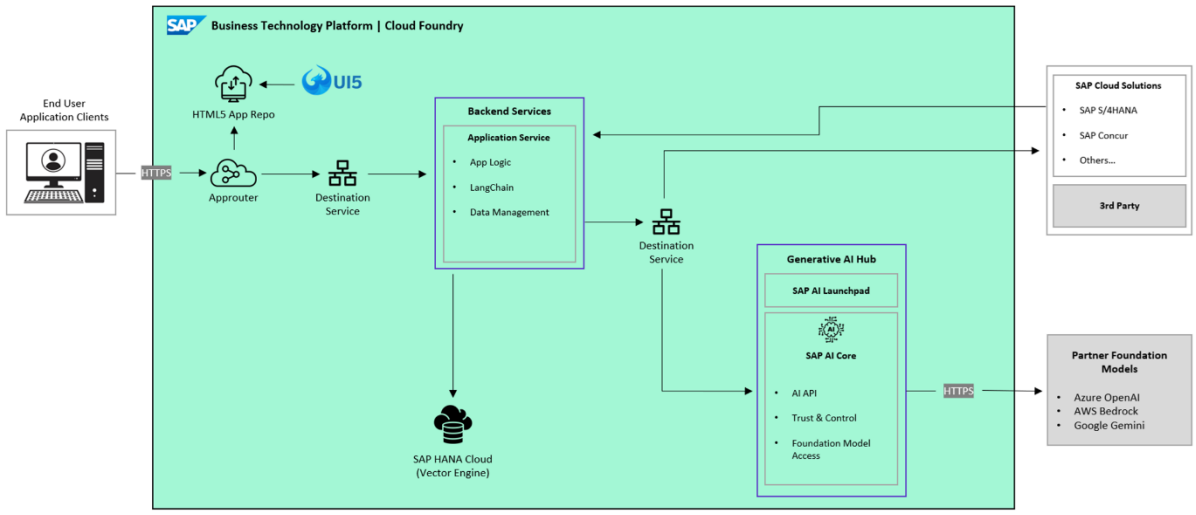


An overview of the process and roles in a contemporary AI project.

[9] https://neptune.ai/blog/ml-engineer-vs-data-scientist

A practical example

The best way to test the value of an AI application, is to build it. To test our own knowledge and skills, we created an AI solution on the **SAP BTP** from scratch using Cloud Application Programming (CAP). We made the decision to build the app in the domain of customer support, as this field can not only benefit greatly from AI support, it also lends itself very well from a functional point of view as it concerns interaction between humans and not machines.



The architecture of the Expertum AI Customer Support application on SAP BTP.

In our use case, we leveraged Generative AI to analyse incoming emails from customers by performing categorization, sentiment analysis, and urgency assessment. This helps customer support in analysing their workload. To make this process more customer or business specific, we also leveraged the aforementioned **RAG**-functionality to store and retrieve past emails and responses to have the model 'learn' from real-world data. We subsequently used the context that RAG provided to create context for a prompt and use that to generate a user-friendly response to incoming emails, making responding to customers easier and more effective. Our use case also broke language barriers by translating the emails and its corresponding responses to a preferred working language (in our case, Dutch) in real-time. Note that you can also connect the application to other SAP Cloud solutions, such as **SAP S/4HANA**, among others. Finally, to add to its versatility, we could also reuse this architecture for other use cases by utilizing Terraform (but that is a story for another time and place!).

In conclusion, we showed a use-case that leverages Generative AI and RAG in order to streamline the workflow of customer support, creating direct value by sending more adequate responses to customers at a faster pace. Important to note is that this application does not simply replace the work of customer support, but rather assists customer support in their workflows, yielding a result that is not only **valuable**, but also ethically **responsible**.

The application we showcased in this section was demonstrated to a **live audience** of over seventy SAP users and customers at the VNSG (Dutch SAP user community) AI and Analytics Day in Amersfoort in November 2024. If you have not seen our application in action yet, but are curious to see what it can do (and how it has evolved further!), feel free to contact us.

Closing remarks

In this White Paper we have touched upon AI as a subject, its history, its place in the SAP portfolio and finally, our vision and first practical applications of it in an (SAP) IT context. We have seen the ups and downs of AI as a research field, explained why it is important to understand the fundamentals underpinning the contemporary technology, and showcased an actual AI application for real-world purposes. At Expertum, we have spent months learning, building, training, testing and failing with AI (applications) to get to the point where we are today: a position from which we can say with confidence that we are able to make AI work in (SAP) IT landscapes.

And so we end this section in the same vein as we finished the preface: We are firm believers of a future where AI can be used effectively, sustainably and responsibly. If you share this vision, would *like to* share it or are simply interested to explore the possibilities you have to leverage Artificial Intelligence in your organisation, please do not hesitate to contact us.