




Microsoft Partner
Azure Expert MSP




Artificial Intelligence and Microsoft Azure AI

Gaining insights from data



Content

Introduction: Artificial Intelligence	4
What is Artificial Intelligence?	6
Key concepts of Artificial Intelligence	6
Examples of use	8
.....	
Machine learning	10
Introduction to machine learning	10
What exactly does a machine learn?	10
How does machine learning work?	12
.....	
Deep learning	16
What is deep learning?	16
The development of deep learning	17
When to use deep learning?	18
.....	
Implementing AI applications with Azure AI	20
Azure AI at a glance	22
Azure machine learning	23
Azure AI services	24
AI builder	28
.....	
SoftwareOne helps you take your first steps into the world of AI	30
.....	



Introduction: Artificial Intelligence


With the success of ChatGPT, more and more companies are discussing, planning and testing the potential applications of artificial intelligence (AI) and machine learning (ML) (see figure). Today, artificial intelligence and data science play a decisive role in gaining insights from growing amounts of data to make informed decisions, improve processes, develop innovative products and gain competitive advantage.

This white paper aims to provide the basic knowledge needed to realistically assess the potential and prospects of current AI technologies, and to outline the first steps towards innovative AI applications. We explain what artificial intelligence is, provide information on technology fundamentals such as machine learning and deep learning, and outline how to implement AI applications with Microsoft Azure AI.

Tip: As a supplement to this white paper, SoftwareOne offers the white paper **“Generative AI: NLP, Language Models and Microsoft Azure OpenAI service”**, which deals specifically with the fundamentals and potential of generative AI – i.e. the AI models behind ChatGPT, DALL-E & Co

Most popular business areas for artificial intelligence

Companies already using or planning to use AI in various areas



Source: Forbes Advisor • Embed

Forbes ADVISOR

AI in US companies: Use and planning (source: [Forbes](#))

What is Artificial Intelligence?

The term “artificial intelligence” refers to systems that can learn and make decisions on their own. Such “intelligent” systems should be able to perform tasks that normally require human intelligence, such as having conversations, translating, assessing situations, driving vehicles or operating machines. For this purpose, AI researchers are developing algorithms and models that enable machines to mimic cognitive functions such as learning, reasoning, problem-solving, perception and decision-making.

The success of AI models depends on often large, high-quality data sets. AI uses data science techniques and methods to collect, clean and process these data sets. Data science is a multidisciplinary field that combines insights from different disciplines such as statistics, mathematics or computer science to gain knowledge from data. In addition, various machine learning techniques are used in AI to train AI models on the data.

Key concepts of Artificial Intelligence

Below we introduce some key concepts of artificial intelligence that are useful for understanding the principles and techniques of AI:

Machine learning (ML)

Machine learning enables systems to recognise patterns and make predictions or decisions without being explicitly programmed. ML involves different approaches, notably supervised learning, unsupervised learning and reinforcement learning (more on this in Chapter 2).

An ML learning algorithm processes training data, looks for structures, patterns and regularities according to a defined objective function, and gradually builds a statistical model whose predictions are tested and improved with further data. After training, the system should be able to handle unknown data.

Traditional machine learning

Unlike deep learning (see below), traditional ML algorithms, such as linear regression, support vector machines or decision trees, require prepared data with a known structure, defined features and only a limited number of feature dimensions for training. Traditional ML algorithms are particularly suitable for small to medium-sized data sets and tasks such as classifying and making certain predictions, where simplicity and interpretability are important. There is also the benefit of having fairly low hardware requirements.

Deep learning

Deep learning is a subset of machine learning that focuses on training deep neural networks. Artificial Neural Networks (ANNs) are computer models inspired by the structure and function of biological neural networks (especially the brain). They consist of interconnected nodes, or “neurons”, organised in layers. “Deep” neural networks have several such layers and are capable of performing complex calculations and processing large amounts of data. In fact, deep learning has enabled significant advances in generative AI, natural language processing (NLP) or the development of predictive models. Read more about deep learning and its pros and cons compared to traditional ML in chapter 3.

Generative AI

Generative AI is a form of artificial intelligence that generates new content based on training data and instructions. AI technologies and algorithms such as Deep Learning, Generative Adversarial Networks (GAN), Variational Autoencoders (VAE) or Transformer architectures are used to generate text including programme code, image, audio and video content, 3D models or virtual worlds. Generative AI differs from discriminative AI, which analyses and evaluates existing data.

Natural language processing (NLP)

NLP or natural language processing is designed to enable computers to “understand” and produce human language. This includes skills such as recognising spoken or written signs and relationships between words, extracting information, summarising or translating text, or responding appropriately to commands or other prompts.



Examples of use

Today, AI offers a wide range of applications in all areas. Companies are increasingly using AI technologies to gain competitive advantage and drive innovation. Examples of common AI applications include:

Chatbots and other natural language applications

Chatbots and other NLP applications are used, for example, to automate customer support and improve customer interaction. Chatbots can handle enquiries, make personalised recommendations or help with bookings. Other NLP applications perform tasks such as text categorisation and generation, text-to-speech / speech-to-text, machine translation or sentiment analysis (see below).

Predictive analytics and forecasts

Predictive analytics uses historical and real-time data to make predictions about future outcomes. Companies use predictive analytics to optimise inventory management, demand and sales forecasting, risk assessment and fraud detection.

Recommendation systems

Recommendation systems are mainly used in e-commerce, media streaming and content platforms. They provide personalised recommendations and improve customer satisfaction by analysing user preferences based on historical data.

Image and video analysis


Computer vision technologies are used in a wide range of industries, including retail, healthcare, manufacturing and security. Image and video analytics support applications such as object recognition and tracking, quality control, facial recognition, autonomous vehicles and surveillance.

Fraud detection and cyber security

AI plays a crucial role in reducing fraud and improving IT security. Machine learning algorithms analyse transaction data, log files and more to identify anomalies, suspicious behaviour patterns and potential security threats in real-time.

Intelligent virtual assistants

Businesses are integrating Intelligent Virtual Assistants (VAs) into their products, services or processes to improve usability or simplify workflows. VAs interact by voice, perform tasks, retrieve information or control smart devices.



Autonomous vehicles

AI is heavily involved in developing autonomous vehicles. AI technologies such as computer vision, sensor fusion and machine learning enable vehicles to perceive their environment, navigate safely and make appropriate decisions.

Supply chain optimisation

AI-powered supply chain optimisation solutions can help companies with inventory management, demand forecasting, logistics planning and route optimisation. They enable more efficient supply chain operations, reduce costs and improve delivery times.

Sentiment analysis and social media monitoring

Sentiment analysis tools analyse social media content, customer reviews and other feedback to understand public opinion, monitor brand reputation or identify trends. With this information, companies can improve their marketing strategies, optimise products and services, and enhance their online presence.

Personalised marketing and customer segmentation

AI enables companies to analyse customer data, behaviour and preferences to create personalised marketing campaigns and targeted advertising. With AI, they can tailor their offers and marketing messages to specific customer segments, improving customer satisfaction and conversion rates.

Machine learning

Introduction to machine learning

A subset of AI, machine learning (ML) develops algorithms and models that can learn on their own to perform tasks for which they have not been specifically programmed, based on input data. ML enables machines to recognise patterns and relationships in large amounts of data and make predictions or decisions based on the processed data.

Probably the most difficult task in computer science is to tell machines exactly what to do. To do this, developers write large software programs, often with thousands or millions of lines of code. Machine learning is an alternative: in contrast to classical programming, in machine learning general instructions on how machines should handle their input data are sufficient for them to find suitable solution algorithms on their own. This has advantages whenever explicit specifications would be too difficult or costly, for example, for complex tasks or large amounts of data, when the optimal algorithm for a task is not known, or when there is uncertainty.

The concept of machine learning can be traced back to the mid-20th century. However, it was not until the availability of powerful computers and large data sets in the 1990s that ML began to take off, becoming arguably the most important AI approach.

What exactly does a machine learn?

Machine learning is a process in which a system derives rules or procedures from a data set to explain the data, i.e. to find structures, similarities or relationships between data points, or to predict future data.

The result of a machine learning system is a model that can be seen as an algorithm for future calculations. The more data the system can process for this, the better it can refine the model and optimise its output. The quality of the training data also determines the success of the model. If the data is flawed, the model's output will be incorrect.


Specific ML tasks differ in the type of prediction or inference required for the problem at hand:

Classification tasks assign data to categories or generally try to answer yes or no questions. A well-known example is the classification of emails into legitimate and spam emails.

Regression tasks attempt to predict (continuous) output values from a set of relevant input characteristics, e.g. property prices based on data such as location, size, amenities, etc. However, many prediction tasks, such as sales or weather forecasting, require more complex models and methods, such as artificial neural networks or the additional use of decision trees.

Clustering tasks group discrete data according to similarity (common attributes). For example, a properly trained model could analyse data from online sales and identify different types of customers.

Anomaly detection, another common ML task, uses statistical methods to identify deviations (outliers) in data. This can be used, for example, to detect attempted fraud or intrusion into an IT network.




What does a machine learn? (Source: SoftwareOne)

How does machine learning work?


The lifecycle of an ML project

Machine learning always starts with the definition of a specific ML task, followed by the provision of appropriate data, such as text documents (e.g. sales reports), numbers (e.g. payment or sensor data), audio data (e.g. voice recordings) or visual data (e.g. photos, medical images or scanned text). The data is processed to be used as training data for the machine learning model of choice. Some of this data is not used for training, but as evaluation data. After several runs and optimisations, the model is able to process new data correctly and can be made available. The life cycle of an ML project also includes the monitoring and management of the productive models.



Lifecycle of an ML project (source: Microsoft)

Machine learning from data can follow different methods (learning paradigms), such as supervised or unsupervised learning (more on this in a moment). Instead of training data sets, machines can be presented with complete learning environments in which to act and gain “experience” (see Reinforcement Learning below).



Supervised machine learning

In supervised ML, models are trained using labelled data sets that associate each input with the desired output. The model learns to generalise these examples and predict the appropriate output for new inputs.

Classification problems are an example of supervised ML. For example, an algorithm can be trained on labelled images of cats and learn to classify images of cats. Regression tasks are also typical examples of supervised learning.

Unsupervised machine learning

Unsupervised machine learning does not use labelled data. Instead, the unsupervised trained model should learn to find patterns in its input data, such as structures, distributions or similarities between input features. It forms an internal representation of the training data, which can then be used for exploratory data analysis or even for generating new data.

A typical task for unsupervised ML is clustering, where the model independently searches for features that can be used to group input patterns. Anomaly detection and dimension reduction (simplification) are also trained unsupervised.


Note: There are also intermediate forms of these types. These include semi-supervised machine learning, where only part of the training data is labelled to reduce the preparation effort, and self-supervised learning, where the required labels are first generated in a separate ML task and then used in further (supervised) learning tasks.

Reinforcement machine learning (reinforcement learning)

Reinforcement machine learning or reinforcement learning does not use training data sets. In psychology, this term describes learning from the consequences of an action in a particular situation or environment. In machine learning, the environment of the learning agent (the learning software system) consists of a set of possible states as well as possible actions that the agent can decide to take. Starting from a random initial state, the learning agent should interact with this environment and, based on feedback (rewards), gradually adjust its strategy to maximise the expected total reward (sum of all upcoming rewards).

An evolution is “Reinforcement Learning from Human Feedback” (RLHF), in which a “reward model” is first trained using human feedback and then used as a reward function. RLHF was one of the methods used in the development of ChatGPT.

Tip: Learn more about Generative AI and GPT model training in our white paper “**Generative AI: NLP, Language Models and Microsoft Azure OpenAI service**”.




Deep learning

Generative AI systems use the method of “deep learning”, which was briefly introduced in Chapter 1. Find out more about their function, history and use.

What is deep learning?

Deep learning is based on “deep” artificial neural networks (ANN), which have many extra layers of neurons between the input and output layers, with each layer processing the input from the previous layer. Each connection between neurons in the network can contribute different amounts to the final output, depending on its “weights”; the weights thus act as adjustable “parameters” of the network.

Layering allows complex processing to be mapped, with layers having different assignments. For example, in face recognition, some layers might respond to individual features of a face, such as eyes or mouth, and another layer might respond to the relative arrangement of these features. However, it is not possible to specify exactly how the input features are processed in the network and how different input signals affect the output. The layers of a deep ANN are therefore also referred to as “hidden layers”.




Traditional Machine Learning vs. Deep Learning (source: SoftwareOne)

The development of deep learning

Concepts for (not yet learning) artificial neural networks or ANNs have been around since the middle of the 20th century. The first learning neural networks were introduced in the 1960s and early 1970s. While these early networks were very limited in scope and complexity compared to modern deep learning architectures, they laid the foundation for the development of more powerful models.

Notable milestones in the development of deep neural networks were Kunihiro Fukushima’s “Neocognitron” in 1979 – the first representative of the now widely used CNN (convolutional neural network) architecture – and the use of back-propagation algorithms for more efficient training of deep neural networks since the 1980s. Another important ANN architecture, recurrent neural networks (RNNs), gained prominence in the 1990s, especially for tasks such as speech recognition, language modelling and machine translation.

In the years after 2000, the improved availability of large data sets (e.g. ImageNet) and increased computing power made training deep networks easier and faster. This has been complemented by architectural innovations such as residual neural networks optimised for large models with hundreds of layers, or attention mechanisms.



When to use deep learning?

Deep learning is achieving remarkable success in many areas, but it is not the best approach for every problem and has its strengths and weaknesses compared to traditional ML methods.

What use cases is deep learning suitable for?

Large and complex data sets: Deep learning is particularly suitable when a large amount of labelled data is available for training. It can automatically learn complex patterns and extract features from data, making it ideal for tasks such as image recognition, natural language processing and speech recognition.

Unstructured data: Deep learning can process unstructured data types such as images, audio and text, making it valuable for image and audio processing or language modelling, for example.

Complex feature structures (feature engineering): Identifying the features in data that are relevant to a task usually requires in-depth expertise. This is where feature engineering comes in: the selection, extraction and transformation of relevant features from raw data to create a representation suitable for ML algorithms. Deep learning reduces the need for manual feature engineering – an advantage especially for complex and abstract features that are difficult to define explicitly.

Advantages of deep learning

High performance in complex tasks: Deep learning models excel in a wide range of areas and especially in complex tasks. They outperform other machine learning methods in tasks such as image classification, object recognition, machine translation and many other language tasks.

Reduced complexity: Deep learning enables end-to-end learning, where the model learns directly from the raw data to produce the desired results. This eliminates the need for manual feature extraction and reduces the overall complexity of the system.

Flexibility and versatility: Deep learning models generate data representations at different levels and can identify relevant features for different tasks, often without being explicitly trained. This makes deep learning extremely flexible and applicable to many problems.

Disadvantages of deep learning

Big data: Deep learning models typically require significant amounts of labelled training data. Obtaining and processing such data can be very time-consuming and expensive. Traditional ML algorithms often perform better on smaller data sets.

Computing resources: Training deep learning models is often resource-intensive and may even require specialised hardware, such as powerful graphics processing units (GPUs) or tensor processing units (TPUs). Traditional ML often runs on low-end hardware as well.


Black box nature: Unlike traditional ML models, complex deep learning models make it difficult to understand their decisions and the learned data representations. This is always problematic when transparency is important – i.e. when people are significantly affected by AI decisions – and motivates the development of “explainable AI” (XAI), which makes AI systems more transparent and understandable.

Implementing AI applications with Azure AI

Companies looking to develop or deploy AI applications face a number of challenges. Training more complex AI models yourself requires sufficient computing resources and expertise. In addition, supervised learning requires the acquisition and labelling of training data. Also, the time required for deep learning should not be underestimated, as several training cycles with step-by-step model optimisation are usually involved.

Those looking to quickly benefit from innovative AI applications should consider a dedicated cloud AI platform. Platforms such as Microsoft Azure offer not only a powerful infrastructure, but also a range of services that facilitate the development and deployment of AI applications.

With Azure AI, there are even extensive pre-trained AI models available a wide range of applications, including ChatGPT and GPT-4 or LLaMA2 from Meta. These models have already been trained for specific tasks with very large amounts of data and usually only need some fine-tuning for individual purposes. Azure AI delivers these models in a manner optimised for enterprise use with Azure's extensive compliance, privacy and security features.




Azure AI at a glance

Azure AI comprises a comprehensive portfolio of cloud services for AI applications – from infrastructure services to pre-trained AI models for various tasks to a versatile application platform for developing customised applications.

Developers can access powerful AI models for image analysis, speech recognition, language processing, and decision making through simple API calls, and train their machine learning models on Azure's high-performance computing (HPC) infrastructure with AI-optimised hardware using familiar tools and open source frameworks such as TensorFlow, PyTorch, or Jupyter Notebooks.

In addition, various AI solutions are available both in Microsoft's SaaS offerings such as Microsoft 365, Microsoft Dynamics 365, and in solutions from Microsoft partners.



Microsoft Azure AI at a glance (source: [Microsoft](#))

Azure AI includes these components, among others:


- **Azure machine learning:** Platform for machine learning
- **Azure AI services:** AI models and tools for cognitive applications and business scenarios
- **AI builder:** AI features without programming via Power platform

Azure machine learning

Azure Machine Learning is a cloud service that enables organisations to manage and accelerate their machine learning projects and processes (ML-Ops).

Machine learning experts and data scientists have tools for many tasks – from labelling data, designing, building, training and deploying models, to visualising and optimising experiments. They can write and run their own code using Jupyter Notebook or train their models without any code at all. They can also adapt a growing number of pre-trained base models, such as open source models or (exclusively) language models from OpenAI. Automated machine learning (AutoML) capabilities automate time-consuming, iterative tasks in the development of a machine learning model.

Developers will also find useful tools for integrating models into their applications or services. Machine learning projects can be easily and securely deployed and scaled, protected by governance, security and compliance features.



Azure machine learning studio: Easy creation of models (source: [Microsoft](#))

Azure AI services

Azure AI services provide out-of-the-box and customisable AI tools and capabilities that enable organisations to modernise and automate processes, gain insights, improve their service, or understand the root cause of anomalies.

With Azure AI services, developers can integrate the cognitive intelligence of powerful pre-trained models into their applications without extensive AI or ML skills.

Azure AI services include both “cognitive” capabilities such as speech analysis (formerly “Azure cognitive services”) and high-level services for specific scenarios (e.g. chatbots or accessibility).

Cognitive services

Basic cognitive services are available via REST APIs or software development kits (SDKs) with client libraries for popular programming languages, and are offered in five main categories:

Image analysis: APIs for image processing and recognition, including customised classifications and face recognition.


Speech: Features related to spoken language, such as speaker recognition, text-to-speech, or transcription and translation of audio.

Language: Language or NLP capabilities for text understanding and analysis, real-time analytics, and building custom question-answer systems.


Decision-making: Anomaly detection, personalisation and content moderation capabilities.

Azure OpenAI service: Access to OpenAI’s powerful language models including the latest GPT-4 and GPT-3.5 Turbo model series, ChatGPT and Codex (via API, SDK or web-based); including content generation, summarisation, semantic search, and natural language to code translation.

Tip: Our white paper, “**Generative AI: NLP, Language Models and Microsoft Azure OpenAI service**”, provides detailed information about natural language processing, modern language models and the solutions offered by OpenAI and Azure OpenAI service.



Azure OpenAI studio (source: [Microsoft](#))



Scenario-specific services

Azure AI services also provides task-specific, high-level services. They combine basic cognitive services with scenario-specific business logic and, depending on the service, can be used via APIs, SDKs, online tools and development environments, the Azure Power Platform (Power Virtual Agents for chatbots) or as Platform as a Service (Cognitive Search).

These out-of-the-box AI services are already optimised for key tasks such as monitoring and diagnosing metrics, extracting knowledge from documents or developing chatbots.

The following services are offered

Azure AI document intelligence: Intelligent processing of a wide range of documents, including document recognition and classification, or the identification and extraction of information (e.g. text, formulas or tables), including predefined templates for contracts, invoices, receipts, tax forms, ID cards or business cards.


Azure metrics advisor: Data monitoring, diagnostics and anomaly detection in time-series data of metrics of various systems, services and resources, e.g. for predictive maintenance of business monitoring applications.

Azure cognitive search: AI cloud search service to identify and exploit relevant information in all types of enterprise content, including semantic search to detect user intent and contextual scoring of search results

Azure immersive reader: An inclusive tool to improve reading comprehension for beginners, language learners and people with learning difficulties such as dyslexia.

Azure Bot service: Quickly create customisable chatbots with pre-built conversation components.

Azure video analyser: AI video analysis to identify and extract information from video and audio files in real time.




Azure Bot Service: Bot framework composer (source: [Microsoft](#))

AI builder

AI Builder provides AI models as a function in the Microsoft Power Platform. The service is directly integrated into Power Apps and Power Automate, where it enables the implementation of AI capabilities for automating processes and extracting insights from data in custom applications.

Even without programming or data science skills, you can create your own models or select and use a pre-built model in five simple steps:

1. **Select AI model type:** according to business requirements
2. **Link data:** select company-specific data
3. **Adapt AI model:** optimise user-defined model (depending on type)
4. **Train AI model:** automatic training process based on company-specific data and specifications
5. **Leverage insights from the AI model:** apply model output on the Power Platform when building your own solutions



AI Builder homepage (source: [Microsoft](#))

SoftwareOne helps you take your first steps into the world of AI


How SoftwareOne supports you

SoftwareOne Data Analytics Kick-Start – agile approach for your AI project

The SoftwareOne Data Analytics Kick-Start makes it easy to get started with Microsoft Azure AI – from requirements engineering and planning, to the technical foundation in the cloud, to your working AI application (prototype, PoC or MVP).

With our proven iterative approach, you can initialise and implement a lean, standardised solution with the most advanced AI services from Microsoft Azure AI and start seeing results quickly.

SoftwareOne Data Analytics Kick-Start



SoftwareOne Data Analytics Kick-Start at a glance (source: SoftwareOne)

CONTACT US TODAY

Find out more at
www.softwareone.com

DE
T. +49 341 2568 000
E. info.de@softwareone.com

AT
T. +43 1 878 10 0
E. info.at@softwareone.com

CH
T. +41 844 44 55 44
E. info.ch@softwareone.com



Copyright © 2023 by SoftwareOne AG, Riedenmatt 4, CH-6370 Stans. All rights reserved. SoftwareOne is a registered trademark of SoftwareOne AG. All other trademarks are the property of their respective owners. SoftwareOne shall not be liable for any error in this document. Liability for damages directly and indirectly associated with the supply or use of this document is excluded as far as legally permissible. © Imagery by: Adobe Stock and Getty Images.