

COGNITIVE SOLUTIONS – A SHORT JOURNEY THROUGH TIME

Artificial intelligence, machine learning, deep learning, cognitive computing ...
these terms are often used synonymously. But what do they mean?
Where and how do we use such solutions today? What are the challenges?
And where is the great potential that's waiting to be released?

By Matthias Loeper

The question as to what intelligence actually is has been intriguing people since the dawn of time but there is still no generally recognized definition. Wikipedia defines intelligence as the ability to perceive or infer information and retain it as knowledge. In other words, intelligence is a measure of how well a person can process information using cognitive functions such as perception, memory, thinking and the use of language. If these functions are carried out by a machine, we refer to artificial intelligence. Experts distinguish between strong and weak artificial intelligence. Weakly intelligent systems are very application related. They use machine learning for clearly defined tasks and process information in a way that appears intelligent. Such systems are used, for example, for weather forecasts, recommendation systems, voice processing and anomaly detection. Highly intelligent systems can think like a person and thus, for example, find out new things and make decisions. The question as to whether strong artificial intelligence can ever exist is still the subject of many a heated debate today.

**IF A MACHINE CARRIES OUT
FUNCTIONS LIKE PERCEPTION,
MEMORY AND THINKING,
WE REFER TO AI.**

The beginnings

The history of artificial intelligence goes back to the invention of the computer in the 1930s. In the early days, computers were mostly used to solve problems that could be described

using mathematical rules but were, for example, too difficult for people to solve because of the sheer enormity of the calculations involved. But the true challenge is in solving tasks which a person finds easy but which cannot be defined as a mathematical rule, such as understanding language, recognizing faces or walking on uneven terrain.

**THE TRUE CHALLENGE
IS IN SOLVING TASKS WHICH
CANNOT BE DEFINED AS
A MATHEMATICAL RULE.**


Machine learning

If knowledge is generated automatically from data, the process is referred to as machine learning (ML). A system learns to recognize patterns and regularities from training data. Machine learning is divided into a number of categories: supervised learning, unsupervised learning and reinforcement learning.

In supervised learning, a system is trained with data sets consisting of input and the expected output. The system thus learns to recognize correlations and make forecasts for future inputs.

In unsupervised learning, the exact form of the result is unknown. The system is fed data and from this data independently generates a model and categories according to which it classifies the data. The aim is to understand the available data, in other words reveal hidden structures and groupings.

In reinforcement learning, the system learns how it should react to situations using feedback. This form of machine learning



Artificial Intelligence (AI): A topic that has been interesting people since the dawn of time.

Artificial intelligence (AI)

1939 – Electro: the humanoid robot Electro, also called the “smoking robot”, was presented at the World Fair in New York in 1939. Electro was controlled using a telephone connection. He could move, count, smoke a cigar and had a limited vocabulary of 700 words with which he could simulate a simple conversation.

1943 – First neural networks.

1948 – Norbert Wiener defined the term cybernetics (the art of controlling) in “Cybernetics or Control and Communication in the Animal and the Machine”.

1950 – Alan Turing aimed to prove that the brain is nothing more than a computer. This was the advent of the Turing test. A test person talks to two interlocutors, one of which is a machine. If the test person cannot distinguish between the human being and the machine, the machine has passed the test.

1950–66 – In the 1950s, the US government put significant funding into projects focusing on machine translation. In the initial stages, sentences were translated word for word before being put together. It was not understood until much later that automatic translation requires extensive knowledge of world events.

1956 – Dartmouth Conference, also known as the big AI bang. The participants came to the conclusion that thinking is also possible without a human brain.

1957/67 – Herbert A. Simon and Allen Newell developed the General Problem Solver (GPS), which was supposed to simulate human thinking. The attempt was abandoned in 1967 but resulted in the development of expert systems.

1965–1975 – First AI winter. Reasons: Recent years were shaped by exaggerated expectations; war in Vietnam.

1966 – ELIZA: Joseph Weizenbaum developed the first well-known chatbot ELIZA that simulated different interlocutors

using scripts. ELIZA became well known for simulating a psychotherapist.

1966 – Machine translation: a report compiled for the US Ministry of Defense came to the conclusion that machine translation is not possible. Research in this area was then effectively shelved for almost 20 years.

1970s – Fight about the ontological status of artificial intelligence. As a result of this argument, weak and strong AI are still seen as contrary positions to one another today.

From the middle of the 1970s – Development of expert systems, e.g. for supporting diagnostic and therapy decisions. But in spite of the considerable investments, they did not meet expectations. Because expert systems cannot learn, the entire knowledge has to be programmed, often with a complex set of rules.

1985–95 – Second AI winter. AI started getting competition when research into neural networks were reactivated. But even this second attempt was too early. There was not sufficient training data and solutions for structuring and modularizing the networks and computers were not powerful enough.

1997 – IBM’s chess computer Deep Blue beat world champion G. Kasparov. At the same time, major investments were being made into the development of robotics.

Around 2010 – The development of the first robots which could optimize their behavior independently with machine learning.

2011 – Jeopardy! Challenge: IBM’s program Watson won against the three top candidates.

2010 to the present day – In around 2010, AI started to be commercialized, particularly in the areas of machine learning and natural language processing. “Deep learning”, the rediscovery of neural networks, played a pivotal role here. Deep learning suddenly made it possible to solve problems that for a long time had seemed impossible to solve. It made the breakthrough in pattern recognition with unstructured data which signified the commercial turning point.

comes closest to human learning. It is suitable for sequential decision processes and enables the automation of sequences which are too complex to be programmed.

Over recent years, machine learning has been given additional impetus by the rediscovery of the significance of hierarchical neural networks, referred to as deep learning. Deep learning is particularly suitable for applications in which large data sets are available from which patterns and models can be derived.

Central: defining an aim precisely

Machine learning provides companies with countless possibilities, from the intelligent automation of processes to the development of disruptive business models. There is one question that is central to every project: What is the aim? In a first step, the current situation, working processes and business processes are analyzed and the aim is defined as precisely as possible. Then a strategy is developed to achieve the necessary transformation. It is not until this point that you can tell whether machine learning can be used sensibly or not and for what purpose.

Creating ML skills

If a company has no experience with machine learning, they should start off with a small project. In virtually every environment there are "low-hanging fruits" in the form of simple application patterns which can be identified fast. This gives those involved the opportunity to have their first experience with machine learning. This is how the business representatives can also start understanding just what you can do with machine learning. Once they combine these insights with their specialist knowledge, it is soon no longer a case of optimizing what already exists but developing entirely new business models. This is where machine learning reveals its full power.

When is ML indispensable?

If all the necessary data is available in digital form, using machine learning would certainly be sensible if large quantities of data or data in an unstructured form (such as e-mails, letters, voice, video, chat, SMS) are available, if the data changes constantly or if the data has to be processed using expert knowledge that cannot be described formally or can only be described with considerable effort. Let's take a closer look at the individual cases:

■ Processing large amounts of data

If you want to make use of data, you have to understand the data in the first place. In other words, you have to be familiar with the internal correlations and be able to relate the data to other data. A person can quickly reach his/her limits when it comes to large quantities of data or great complexity. And this is when machine learning can help.

■ Automating processes with unstructured data

If a company receives 20 inquiries a day, a person can process one after the other. If, on the other hand, the company receives 20,000 inquiries a day, they have to be automated. This is where machine learning can be used, from triage through to the fully automated answering of inquiries.

■ Processing constantly changing data

When processing data whose patterns or internal correlations change quickly, it can be difficult to use classic, fully programmed systems of rules because these would have to be adapted constantly. A neural system can be of help in this case.

■ Gathering expert knowledge that cannot be described formally

A further area of implementation for machine learning is for tasks in which the experts cannot describe exactly how they solve the task, e.g. because a lot of experience is necessary to carry out the task or because the process is something that the person no longer gives any conscious thought to. If a person is being trained for a new job, he/she usually learns from examples. Machine learning can also be used in such cases. But the machine requires much larger amounts of training data than people do.

**IF A COMPANY HAS
NO EXPERIENCE WITH
MACHINE LEARNING,
THEY SHOULD START OFF
WITH A SMALL PROJECT.**

Where is ML used today?

Machine learning is already used today in a number of areas. Familiar areas of use are:

- Forecasts, predictions: product marketing and maintenance, weather reports, predictive policing
- Voice processing: voice recognition, natural language processing (NLP), language synthesis
- Classification and structuring of data
- Recommendation systems: Amazon, Spotify, Netflix
- Personal assistants: Siri, Amazon Echo, Google Home
- Chatbots: Service Bot Swisscom
- Anomaly detection: discovery of credit card fraud, monitoring, alerting
- Medical diagnostics, epidemiology and biometrics
- Image processing and pattern recognition
- Robotics: movement control, sight, etc.

In spite of the large range of application areas, there are areas which currently cannot be covered by machine learning because they are too complex.

Obstacles when integrating ML

The effort involved in integrating machine learning depends on the quality, the amount and the composition of the available data. Often, data from legacy systems has to be incorporated, access to required information is not possible or difficult, or data cannot be correlated for reasons of data protection. With too great a data volume it could be that the solution is not scaled; with too low a volume, machine learning could well be impossible. Or there is not sufficient expert knowledge for the extraction of the required information from a system.

TO BE ABLE TO MAKE
A COMPUTER THAT THINKS LIKE
A PERSON, WE WOULD FIRST HAVE
TO UNDERSTAND HOW THE HUMAN
BRAIN WORKS, I.E. TO “CRACK
THE BRAIN CODE”.

Dealing with errors and evaluations

Another obstacle to using machine learning is the lack of predictability or verifiability. Conventional computer programs behave deterministically, in other words, the rules according to which they make their decisions are comprehensible and can be verified. When we teach systems to learn like people, there are certain disadvantages. Human beings do not work on a digital basis and their behavior is not always predictable. That is also the case with machine learning, particularly when deep learning is used. A program's behavior and the basis for its decisions which occur in the process are generally not totally comprehensible. It is often the case that we do not know how they came up with a certain result.

Another aspect is error handling. Systems using machine learning also make mistakes. Sometimes the learning algorithms use false assumptions (distortion) or react too sensitively to variations in the data (variance). Therefore, the systems have to be designed so they can cope with errors.

Implicit judgments are also tricky. If a model is trained with data containing implicit judgments, for example, that are not gender-neutral or are racist, the model also learns these judgments. The dangerous aspect of this is that judgments in data can under certain circumstances be difficult to detect. The less we understand the data, the more tricky the situation. The systems learn mistakes and because human beings do not recognize the mistakes, they cannot intervene to put things right.

Where does ML reach its limits?

As the article is coming to its close, let's be brave and take a look into the future: Have we now finally made the great break-

through in our dealings with artificial intelligence and will it now increase exponentially, only held back by the performance of the computer? Or have we indeed taken a large step but have now once more reached a plateau?

To be able to make a computer that thinks like a person, we would first have to understand how the human brain works. We would have to “crack the brain code” as Pascal Kaufmann, CEO of StarMind, puts it. Jeff Hawkins is working in a similar direction with the company Numenta and what is referred to as Hierarchical Temporal Memory (HTM). This is a form of hierarchical neural networks which focus on the temporal aspect and try to reproduce fundamental mechanisms of the neocortex.

This is something that is also being pursued by the Human Brain Project (HBP), a major project of the European Commission. The project was initiated by Henry Markram, Professor at the ETH Lausanne, and originally intended to build a computer model of the human brain by 2023. So far they have only managed to reproduce a tiny piece of the cerebral cortex of a rodent. But if the researchers were able to discover key regularities when building the model, the original aim could once again become the focus.

While research on strong artificial intelligence is still in its infancy, enormous progress has without doubt been made in recent years in the area of machine learning with the rediscovery of the significance of neural networks. Today, virtually every individual task, in isolation, can be completed more efficiently by a machine than by a human being. Machine learning also enables the introduction of new business models and is having a disruptive effect in many areas. As our examples show, machine learning is already being used in many areas today. But there is still a lot of untapped potential and a lot to do until we have fully exploited all the technical possibilities we have available to us today. ■

Matthias Loepfe

Matthias Loepfe, who has a degree in electrical engineering, significantly shaped the development of AdNovum in the early years as Technical Lead and later CTO and co-owner. In 2003 he sold his share and then focused on cyber crime investigation and digital forensics for over ten years. In 2016 he returned to AdNovum and now as Head of AdNovum Incubator and together with his team is researching the suitability for daily use and disruptive strength of innovative technology. Not possible? No such thing. Matthias Loepfe is passionately dedicated to finding the most elegant solution. He proved his skills in craftsmanship when he was younger by rebuilding VW buses. Otherwise he likes being in the great outdoors.



Matthias Loepte: Explores new technologies.