

Learning Machines

Artificial intelligence and machine learning are becoming one of the most important tools of the future. “Machine learning will change our world, just as the internet and computers have done,” says a convinced Robert Legenstein from TU Graz.

Birgit Baustädter

A machine that makes decisions as a matter of course. Which can talk, play and give tuition in a meaningful way. Or: a swarm of heavily armed drones that wants to subjugate humanity. This or something similar is the common image of artificial intelligence. But the reality (still) looks a little different.

“Artificial intelligence has always been understood to mean what was not yet possible by machine at the time,” says TU Graz researcher Robert Legenstein, attempting a definition. “Fifty years ago, that was playing chess, for example. Then it was image and speech recognition. Today, all this can be done by machines.” All of these capabilities are now part of what is known as “weak artificial intelligence”. It is “strong artificial intelligence” that is currently still in the stars – in other words, cognitive systems that have human-like abilities. “Today’s artificial intelligence lacks the ability to generalize. It is perfectly trained for a certain task, but cannot transfer this knowledge to other situations on its own.” For example, a properly trained algorithm can play chess beautifully, or even the more complex game of Go. But even in ludo it wouldn’t stand a chance without prior training.



Robert Legenstein deals with biologically inspired artificial intelligence.

Lunghammer – TU Graz



GRAML – short for Graz Research Center for Machine Learning – is currently being set up at TU Graz. According to its director, Robert Legenstein, it is intended to create a TU Graz-wide network linking researchers who are working on machine learning in different ways. “The initiative comes from institutes in the field of computer science. But we’ve also deliberately brought colleagues from physics and chemistry on board, for example, because we see a lot of potential in these areas.”

SPIKING NEURONS

In his work, Robert Legenstein, from the Institute of Theoretical Computer Science, takes the currently most powerful and yet most energy-efficient computer as his model: the human brain. In a nutshell, he deals with biology-inspired artificial intelligence and builds neural networks – mathematical structures that are similar to the human brain. “The brain is so energy-efficient because only the neurons that are currently needed to pass on >



the information are active. All the others are 'asleep' and don't consume any energy. These are called 'spiking neurons'." Using this way of working, Legenstein and his team build spiking neural networks that can complete complex computational tasks with little energy.

Most recently, as part of an international team, Legenstein developed a robotic arm controlled by a neural network and shaped like an elephant's trunk. "The control problem is very complex because 300 individual motors have to be controlled in coordination with each other to enable smooth movement. This control task can be done by our neural network and is additionally energy-efficient."

UNCERTAIN COMPUTERS

Robert Peharz, from the Institute of Theoretical Computer Science, is also working on foundations of artificial intelligence and machine learning. His main interest lies in probabilistic machine learning, that is, he incorporates probability theory into his algorithms. "Probability theory already covers a fair bit of what artificial intelligence is supposed to do," he explains. "Most importantly, it takes uncertainty into account, represents dependencies in the data, and comes with an integrated reasoning process." In this way, probability theory allows the possibility of learning from data and updating to new situations. "One of my favorite examples is when I ask the question: Can Tim fly? When we assume that Tim is human, then the probability for this event is rather low. However, if I add the information that Tim is a bird, then the probability increases dramatically. But, if I then add that Tim is sick, the probability goes down again. This is an example of non-monotonous reasoning, which is present both in probability and everyday human reasoning."

The most important aspect of machine learning as a tool of the future is its ability to adapt. "Essentially, I can write a meta-program for which I provide the rules of some generic task – for example, 'distinguishing between different things'. This meta-system can then be adapted to more specific tasks – for example, in waste-sorting units or quality control."

According to Robert Peharz, the future journey of AI is uncertain, as it is a scientific field which is simultaneously inventing and researching itself. "To quote Niels Bohr: Predictions are hard, especially if they are about the future."




Franz Pernkopf.

Lunghammer – TU Graz

**A sleep mode
for Siri and Alexa.**

HEARING MACHINES

For Franz Pernkopf, Institute of Signal Processing and Speech Communication, artificial intelligence and machine learning is one thing above all: a tool. "Especially a tool for lazy people," the researcher laughs. "We are faced with complex problems that can no longer be described using easy mathematical models. That's why we use learning algorithms that look for correlations from large amounts of data that would take me, as a researcher, vast amounts of time." Especially in dynamic situations, artificial intelligence has enormous advantages – for example, when it comes to Pernkopf's field of research: language. Hearing machines have now arrived in all our lives in the form of voice assistance systems. But, proper machine hearing is not trivial, as Pernkopf explains: "Due to different ways of speaking, dialects and, above all, background noises, it is hard to assign a particular word to a single acoustic wave. But I can train an algorithm with a large database." This can be taken so far that individual speakers can be filtered out of, for example, a cocktail party situation.



Thomas Pock improves the vision of machines.

Lunghammer – TU Graz

What is machine learning?

SEEING MACHINES

Thomas Pock, Institute of Computer Graphics and Vision, works in a similar field, but his focus is on a different human sense: vision. Our visual cortex can capture images and recognize objects in a fraction of a second, even if they are barely visible or only fragmentary. Thomas Pock bases his algorithms on the way the visual cortex works in order to automatically optimize images in medical diagnostic procedures, for example. As application-oriented as his research sounds, it is nevertheless very much in the area of basic research. “I am actually a very big advocate of ‘handmade’ mathematical models, but I see that they are far from sufficient to describe reality,” he explains. “So I take algorithms that I have designed by hand and give them some free parameters. Then I get them to learn from existing data and improve the model. If I add more and more parameters, I eventually end up with deep learning – a relatively simple but heavily parameterized model in which we can no longer under-

stand at all how the algorithm comes to its conclusions because we can’t interpret the learned parameters.”

Pock thinks the vision of an actual thinking system is still in the distant future – possibly not even realizable at all. “Artificial intelligence is currently to a large extent pattern recognition, not real intelligence. But it’s a great tool when researchers get stuck at certain points and the algorithm can still pick things out of the data because it’s much faster and of course it doesn’t get tired.” >

**Computer Vision:
TU Graz Researchers Define
New State of the Art.**



Olga Saukh.
Lunghammer – TU Graz

MACHINE LEARNING – THE TECHNOLOGY OF THE FUTURE

Machine Learning is a method of Artificial Intelligence and just as old – both had their beginnings in 1956 at the Dartmouth Conference, which dealt with Self-Improvement among other things. A lot has happened since then.

Machine learning emerged as a method from the 2000s onwards. In particular with the boom in deep neural networks, which moved into the centre of scientific interest in 2012 with the paper “ImageNet classification with deep convolutional neural networks” and have since experienced a real boom. I would even go so far as to say that machine learning is the technology that will only allow us in the future to make sense of the huge amounts of data that digitalization is generating.



Horst Bischof
Vice Rector for Research
Oliver Wolf

FLOWING AIR

Olga Saukh researches complex systems for both the Institute of Technical Informatics at TU Graz and the Complexity Science Hub Vienna. Opinion formation in societies, for example, is a complex system. But also is, for example, the movement of fine dust particles in the air, which Saukh has been studying for several years. The focus is on air quality forecasts, which should enable legislation to react flexibly to current changes. “We include the air currents that can blow particulate matter from one region to another, for example, and thus cause short-term changes in air quality,” says Saukh, explaining her approach. In previous models, these air currents would have had little to no influence. “We use machine learning to improve these short-term predic-

tions. In this area, artificial intelligence is far superior to expert models. But when it comes to long-term predictions, the expert models are still more reliable.”


In addition to air pollution, Saukh is also working on the combination of embedded systems and machine learning. “If I want to improve current models, I make them bigger and expand the data sets. Which of course goes hand in hand with higher energy requirements and is not feasible on small systems such as smart watches. So I’m trying to find small, sparse but powerful models and optimize their execution to the underlying hardware capabilities and resource constraints, so that the overall system can operate for months or years on a single pair of batteries. Small is mighty is our motto.”

Machine learning is now used in virtually all areas of science and achieves excellent results. The best-known example is probably the Alpha Go algorithm, which was able to defeat the reigning world champion, Lee Sedol, in the complex Japanese game of Go in 2016 thanks to deep reinforcement learning. AlphaFold, a system also developed by DeepMind, can predict protein structures based on the amino acid sequence of the protein and achieved top scores in the Critical Assessment of Techniques for Protein Structure Prediction (CASP) competition in 2018 and 2020. And machine learning algorithms are also driving development in the field of autonomous driving. The reason for the rapid progress is probably also the fact that many systems are open source and can thus be viewed and further developed by anyone.

Despite all its successes, machine learning is still a very young field. Far-reaching developments are still needed

in both theory (we often don't understand why an algorithm works) and method (algorithms need too much data and too much energy). But I dare to predict that no data-based science will be able to do without machine learning in the future. Only in this way are we able to make sense of the data and create forecasts. Some researchers even go so far as to say that machine learning will change science itself. Algorithms, they predict, will formulate hypotheses in the future, and we researchers will design experiments to prove causality – something algorithms cannot do at present.

Regardless of which scenario occurs, machine learning will play a central role and is one of the most important technologies of the future. ■



RACING NANOPARTICLES

Solid-state physicist Oliver Hofmann uses a playful approach in his research on machine learning and wants to relieve researchers of tedious routine work. “We want to have learning algorithms search for promising material structures that our team would otherwise have to spend weeks painstakingly digging through,” explains Hofmann. But the algorithm does not decide on its own: “We then look at the result and of course recalculate the most attractive structures ourselves. The algorithm can recognize correlations based on previously trained knowledge, but not whether they are also scientifically correct.”

Hofmann would then like to have the structures found in this way replicated and tested in an experiment. “But it’s also a lengthy job to build that structure and often we as researchers have no idea how to even go about it.” Here, too, a trained algorithm should give support and learn to build structures independently. Currently, however, this technology is being explored in an enter-

taining setting: with a “car race” on a nano scale. In the process, nanoparticles – the cars – have to be steered through a course by means of an electrical impulse under the STM – the scanning tunnelling microscope. “All of this takes place in the quantum realm. That means these processes are not deterministic all the time – even if I set up the same electrical impulse, the particle will not move the same way all the time.” The world’s best “racers” would currently manage to actually move the particle every second time they tried. “Whether it then also moves in the right direction is another question,” smiles Hofmann. His team has now trained an algorithm instead of a human “driver” and has already achieved great results: “A human has to have years of experience in order to move the particle efficiently. Our algorithm only needed a few weeks to manoeuvre through the course at least equally fast. That’s a success.”

THINKING MACHINES

The possible applications of artificial intelligence and learning algorithms are therefore diverse – especially as a tool in research. And to return to the vision at the beginning – the thinking robots. Robert Legenstein has a clear opinion here: “I think it is possible up to a certain extent. But they will not behave exactly like humans. Humans are social, sentient beings and a robot does not grow up in the same environment and thus does not have the same experiences. But he might well develop certain human-like behaviours in an artificial environment.” ■