

Digital Jan 3, 2019

The most important technology for the future

Rapid technological progress in the field is making artificial intelligence (AI) one of the key technologies for the 21st century – thanks to Big Data and exponentially increasing computer capacities. But how does it actually work?



It is 50 years since AI became a mainstream concept for the first time – in the form of a Hollywood film. In 1968, director Stanley Kubrick had the HAL 9000 supercomputer take control of the spaceship in his science fiction epic “2001: A Space Odyssey”. A machine that’s more intelligent than humans?

As an academic discipline, AI was 12 years old when the film was released, having arrived at renowned Dartmouth College, in New Hampshire USA, in July 1956. That’s when a group of ambitious mathematicians and electrical engineers convened at the Dartmouth Summer Research Project on Artificial Intelligence, a project was initiated by John McCarthy, who invented LISP – the world’s second-oldest programming language.

Birth of “artificial intelligence”

After an industrious few weeks that summer, the ten invited thinkers had produced reams of dense writing and many ideas. Talking machines; networks based on human brains; self-optimising computers; and even machine creativity seemed to be within reach of this euphoric founding generation. Their most important development though, was the term “artificial intelligence”, the coining of which created a new discipline that would fascinate people worldwide from that moment on – and in fact it caught on quicker than anyone expected.

The very same year, Arthur Lee Samuel – one of the participants in the conference and a computer scientist at Massachusetts Institute of Technology (MIT) – taught an IBM 701 computer how to play the board game checkers. His program used a method whereby the machine could learn from its own experience, particularly in later versions. In 1961, it played the Connecticut state champion – and won. This approach represented the basic idea of AI in action: software learning on the basis of large quantities of data.

Singing computer

Also in that year, a type 704 computer learned the song “Daisy Bell” at Bell Laboratories, and reproduced it using speech synthesis. This evidently appealed to Stanley Kubrick, as he had the HAL 9000 supercomputer sing the same song in his film. To the masses at the time, all of this was pure science fiction; but today, no one falls off their chair in surprise if their computer plays music. It’s another of HAL 9000’s abilities that remains further out of reach: “strong” or “general” AI, meaning AI that comprehensively imitates or could even replace humans, remains a utopian dream.

The Turing test is applied to determine whether an AI development is on a par with humans, and while no technical system is set to pass the test in the foreseeable future, there are some things that machines are already able to do better than people. For example, they are fantastically useful in analysing large quantities of text or data, and they form the bedrock of internet search engines. Embedded in countless smartphone apps, we carry this “weak” AI in our pockets everywhere we go – and as users, we are mostly hardly even aware of it. But anyone who talks to Alexa or Siri is also having their sentences analysed by AI algorithms; John McCarthy made a dry comment on the fate of AI applications: “As soon as it works, no one calls it AI any more.”

Deep Blue beats Chess World Champion Garry Kasparov

He had a point; but before that happens, there’s widespread amazement every time AI passes another milestone, for example in 1997 when chess computer Deep Blue beat Chess World Champion Garry Kasparov. Games are always a popular testbed for AI scientists, and they also offer good opportunities for publicity.

One example is a TV game show called Jeopardy, which involves candidates having to identify the right question to which a given term is the correct answer. The set tasks were generally worded to be deliberately ambiguous, and to require the linking of several facts to find the right answer – making the challenge much more difficult. However, the “Watson” IBM system managed to beat the two human record holders in 2011, after being fed with 100 gigabytes of text. Rather than relying on an individual algorithm, Watson simultaneously used hundreds of them to find a potentially correct answer via a path. The more algorithms independently reached the same answer, the greater the probability that Watson had come to the right conclusion.

DeepMind beats “Go” world champion Lee Sedol

The next bit of excitement came from DeepMind, a London based start-up that was founded in 2010 and integrated into the Google Group in 2014. It developed an AI application that optimises itself when learning games. AlphaGo set itself the target of beating a human “Go” world champion – which was considered an almost insurmountable task given the extreme complexity of this strategy game. AlphaGo achieved its target for the first time in 2016, defeating the reigning world champion Lee Sedol from South Korea: a long-awaited milestone. Currently, the AlphaZero program only defeats itself – because it foregoes human sample games and instead learns only from playing on its own: human players no longer have any chance of winning against AlphaZero.

This feat is made possible by artificial neural networks. Neurons are nerve cells which form a network to which an individual task is allocated, such as vision. An apparently endless number of neurons are dynamically connected within the human nervous system. The human brain learns by adjusting the density of these networks on an ongoing basis; paths that are frequently used become stronger, while neglected connections waste away.

Artificial neural network

An artificial neural network tries to replicate this structure. Artificial neurons networked together take in input values and feed this information into neurons created in lower-level layers. At the end of the chain, a layer of output neurons delivers a result value. The variable weighting of the individual connections gives the network one particularly notable property: the capacity to learn. Today, networks are increasingly based on these levels; they are more complex, and further interlaced – that is, deeper – thanks to increased computer capacity. Some deep neural networks are made up of more than 100 of these series-connected program layers.

However, AI has to be trained – in a process also known as Deep Learning. In this process, systems receive corrective feedback from an external source, for example a human or another piece of software. The system makes its conclusions from the feedback it receives – and it learns.

Promising practical tests

Porsche CIO Matthias Ulbrich believes that AI is the most important technology for the future, and that it will help us to dedicate our time to the things that really matter. “AI will play a part in value creation. In the same way that robots already take the physical strain off us today, AI will support us in thinking and decision-making during routine work,” he explains. The development departments have a lot of work to do before we reach that point. One key consideration in this work involves the aspects of security and personal privacy.

At Porsche, the subject has been taken up by Tobias Große-Puppendahl and Jan Feiling from the main Electrics/Electronics Development department. Developments such as personalisation, swarm intelligence, and the protection of the private sphere all require AI in order to preserve overarching privacy when collecting and exchanging data. The team aims to minimise data exchange by using “federated learning” wherein a local AI system inside the car learns from the user’s behaviour. For example if a driver says, “I’m cold,” the AI should turn up the heating. It passes its learning success – or to put it another way, its experience – on to the cloud and the global AI instruments installed there, while specific data such as language protocols can remain in the car. Ultimately what is key is the intention behind the data: that each user expresses a wish in their own way, but expects the same result. Think of meeting a person whose language we do not understand, but they are able to make it clear if they’re feeling cold.

Pure science fiction

Of course, HAL 9000 from Stanley Kubrick's 2001: A Space Odyssey is also capable of doing that. But an AI rebellion against humans is pure science fiction – at least for the moment – and teleportation as seen in Star Trek will probably forever remain a utopian dream. After all, good science fiction doesn't exclusively reflect real cutting-edge technology that is little-known to the general public – such as the singing computer – but also explores the realms of incredible fantasy. Dresden-based AI specialist Professor Sebastian Rudolph believes that future machine rebellion scenarios are extremely far-fetched given the current state of technology. He says that, as is the case with all tech, AI could be misused – and in fact that mistakes could be made in its implementation.

So perhaps we shouldn't be any more or less afraid of this type of development than we are of technical progress in general. And looked at this way, it makes sense for all of us to participate in shaping this technical progress ourselves. That's what Tobias Große-Puppendahl and Jan Feiling have internalised at Porsche – and in fact in the best of the company's tradition, following Ferry Porsche himself: "We couldn't find AI that appealed to us. So we built it ourselves."

Link Collection

Link to this article

<https://goo.gl/rjzxeu>

Media Package

<https://newsroom.porsche.com/media-package/porsche-artificial-intelligence-technological-progress-big-data>