


Business is waking up to the idea of deep learning

By [Toby Walsh](#)

15 Jun 2016

In the movie *Transcendence*, Johnny Depp plays Dr Will Caster, a researcher in artificial intelligence at Berkeley trying to build a sentient computer. [Stuart Russell](#) is Will Caster's real life equivalent. He works on artificial intelligence at the University of California at Berkeley, and is co-author of [the definitive textbook on AI](#). He has also been very vocal about the risks of research in AI succeeding.



©Nenad Cerovic via [123RF](#)

Earlier this year, Google's DeepMind taught a computer program to play a wide variety of Atari video games at a superhuman level in a matter of hours. The program was given no background knowledge. It learnt every game from scratch. Stuart observed that "If your newborn baby did that you would think it was possessed." Should we be impressed or distressed by such progress?

The secret sauce

The secret behind DeepMind's success is deep learning, an exciting and very active branch of machine learning. Machine learning is driving the next revolution in computing. Instead of programming a computer laboriously by hand to do a task, we simply get the computer to learn to do the task.

[Games are a particularly good domain in which to apply machine learning](#). There are very precise rules to follow. It's easy to work out who wins. And the computer can bootstrap its performance by playing itself millions of times.

Deep learning is, however, good at much more than learning how to play games. It is transforming how computers transcribe speech into text, recognise images, rank search results, and perform many other tasks that require intelligence. Deep learning uses a 'deep' neural network, loosely modelled on the human brain. It's deep because it has half a dozen or so of layers.

These layers are critical to success. They permit the neural network to pick out features. For example, in recognising images, the intermediate layers recognise features like edges and corners. Also critical to success is lots of data and computing power. Deep learning needs a lot of examples from which to learn. And the learning itself is often run on multiple specialised graphical processing units (GPUs).

What's it not good for?

Deep learning looks set then to be an important component in the AI toolbox. However, it is unlikely to solve all of AI. Deep learning appears to be especially good at perceptive tasks like processing speech or images. However, it is more challenged with tasks that require higher level or strategic reasoning. Think planning how to build a factory. Or solving a complex mathematical problem.

DeepMind's program learnt to play simple reactive games like Pong and Space Invaders. But it didn't do well at PacMan as this requires planning ahead to deal with the ghosts.

There are some other challenges. For instance, deep learning requires lots of data. Humans, by comparison, can learn from just a few examples. There are many domains where learning is difficult, painful or even deadly. We will therefore also need methods that are quick to learn. Deep learning is also largely a black box. There are many domains where we will want a computer to explain its reasoning or to justify its answer.

Finally, there are many domains where we will want guarantees. The air traffic control program should never let two aeroplanes into the same airspace. The autonomous car should always stop at a red light. Deep learning comes with no such guarantees.

Nevertheless, there are many businesses that are waking up to the idea of building new services powered by deep learning. It is sure to play a critical role in driving autonomous cars, ranking search results, recommending products, identifying spam email, trading stocks, and interpreting medical images.

How will this all end?

Deep learning is still a long way from matching our intelligence. Deep learning uses thousands of neurons, and millions of connections. But the human brain has billions of neurons and trillions of connections. It remains a major scientific and engineering challenge scaling to the size of the human brain. After all, the human brain is the most complex system we know of in the whole universe.

And then there are a bunch of difficult problems like consciousness and emotions. These seem intimately connected to our intelligence, and we have yet to understand or replicate them in silicon. So we are some way off from Transcendence. The Technological Singularity. The moment when computers start to improve themselves and their intelligence snowballs, soon far exceeding human limits.

[Many researchers in AI actually doubt whether we'll ever get to this moment.](#) That we will run into scientific, engineering or other road blocks which will prevent any run away in machine intelligence.

This won't stop us building machines that will transform our lives. And we need to start planning for that future now, preparing for the time when many jobs have been automated.

ABOUT THE AUTHOR

[[<https://theconversation.com/profiles/toby-walsh-51> Toby Walsh]] is a Professor of Artificial Intelligence at UNSW, and Research Leader at Data61 (formerly NICTA) in the Optimisation Research Group where he leads the Algorithmic Decision Theory project. Data61 is Australia's Centre of Excellence for ICT Research.

For more, visit: <https://www.bizcommunity.com>