

THE RISE OF THE MACHINES

by Joyrene Thomas

Move over Big Data. There are two new buzz words in town: ‘artificial intelligence’ and ‘machine learning’. The claims being made for these technologies are just as significant, if not more so than those made for Big Data three to five years ago. So, how widespread is the use of artificial intelligence today? And what are the implications for the payments industry?

The rise of powerful artificial intelligence will either be the best or the worst thing ever to happen to humanity, according to Professor Stephen Hawking. Meanwhile the chief executives of Microsoft and Google have said that the technology will change not only computing but every industry and business process. Artificial intelligence is trending high in the hype cycle. It seems as though every company has a product claiming to have artificial intelligence built in, but that is not quite how it works.

Artificial intelligence, or AI, is essentially making computers do what minds do. The term was originally coined in 1956 by John McCarthy, an American computer and cognitive scientist, to name what had previously been called ‘computer simulation’. As an academic discipline, AI encompasses various technological and scientific areas. In her book *AI: Its Nature and Future*, Margaret Boden divides AI into five major types, each including many variations. For completeness, these are classical or symbolic AI. This is sometimes charmingly known as GOFAI (good old-fashioned AI). Secondly, artificial neural

networks or connectionism. Additionally, evolutionary programming, cellular automata and dynamical systems.

OK COMPUTER

Many claims around AI are definitely artificial. AI is used to mean automation or ‘doing things with data’. Sometimes even the automation is artificial. A human or manual process is masquerading behind the computer facade. “There’s almost nobody in the world actually doing artificial intelligence – only very few people,” says Martin Sweeney, co-founder and CEO of fraud prevention firm Ravelin. Some of the deep learning aspects of Deep Mind, which is Google’s technology, and some of what Facebook and Amazon are doing could be classified as AI. “What you are seeing is the domination of artificial intelligence among one or two big players, who are hoovering up talent from academia. No-one else can afford this. Actually, not many people need to,” continues Sweeney.

“You still get fantastic results using more conventional technologies with very many

Many claims are ‘artificial’ artificial intelligence – automation or ‘doing things with data’.

similarities with what we are calling artificial intelligence. I think that ‘machine learning’ is a much more accurate description of what people are actually doing, which is to say statistics with reinforcement.”

A company could take a simple pass through a data set to determine what makes customers good or bad. This is supervised, feature extraction-driven statistics. “The machine learning part comes when the world changes and I need to re-think my strategy. This is about taking your data feed and re-generating your models as quickly as possible to re-encapsulate your understanding of the problem you are solving,” says Sweeney.

artificial intelligence

LESS SCOPE, MORE SPEED

Computers are intelligent — at least they can be programmed to be. There are a couple of caveats, though. The intelligence is niche, and sometimes the methods used are very unlike human minds — and perhaps all the better for it.

“One of the criticisms of the term [artificial intelligence] is that it made people think that what was going on behind was that computers would act as intelligently as a human being. That is still far down the road,” says **Dr Ulrich Dorndorf, chief technology officer at advanced optimisation software company INFORM.**

“When computers are intelligent, it works well when they are intelligent in certain niche areas. We have taught computers to play chess and they are successful to the degree that no human can now compete with any of the good chess programmes, but it’s a niche intelligence.”

The practical applications have come by harnessing this niche intelligence. Also by programming computers to think and learn differently, rather than replicate how humans would perform the task. The pattern recognition aspects of machine learning are widespread and already part of our lives. It helps Google serve adverts on websites, depending on what the user has looked at in the past. It is part of e-mail spam filters. Naturally, it is also used by banks in transaction monitoring, credit scoring and risk management.

Machine learning techniques enable companies to analyse huge amounts of data — and to do so quickly. “To pinpoint fraud in real-time, you have to adapt to constantly new patterns. That’s the main complexity: too much data has to be combined that it is close to impossible for humans to react in the time period needed,” explains Konrad Hochmuth, strategic business developer, INFORM.

Whereas in the past a bank may have had two-three days to cross-check a payment, real-time payments are typically on the beneficiary’s account within ten seconds. Real-time payments is driving real-time decisioning. The regulators are in turn driving the banks in this direction. The European Central Bank is expecting “at least one instant payment solution in euro” to be available “to all payment service providers in Europe” towards the end of 2017, according to its website. Irrespective of whether it is regulator-pull, customer-push or a

combination of both, numerous other markets are also exploring or implementing real-time payments. The payments industry crossed the rubicon long ago. Machine learning is part of the present and will be part of the future.

BEYOND THE BLACK BOX

Users and suppliers of machine learning technology will have to grapple with various questions, some of which are not new. One of these is causation. Humans like to make sense of the world through causal links. We like to believe that every effect has a cause and vice versa. So, how do we explain the workings of machines to humans?

There are techniques in artificial intelligence and machine learning that are more prone to be understood by humans, explains Dorndorf from INFORM. For example, with rule-based systems where experts represent their knowledge in the form of rules, system operators can look up which rules led to a conclusion. Conversely there are techniques that will essentially remain a ‘black box’, at least for the moment. “Neural networks are an example of a ‘black box’ approach. People and scientists recognise that they work quite well in certain areas, but even from a scientific point-of-view the question is still open why this is so,” says Dorndorf. Naturally, it is possible to layer the various technologies. “We find that our customers, as users of the technology that we deliver, prefer approaches where the logic does not remain a completely ‘black box’.”

Hochmuth makes an important distinction between humans in the discussion about causation. “The human specialists in the bank are very much interested in understanding

“It works well when computers are intelligent in certain niche areas.”

Dr Ulrich Dorndorf, INFORM

why transactions have been declined. This is their number one resource to optimise the system, and see future trends that may not even be in the data. These humans are hired

It may seem counterintuitive but having less information can lead to better predictions.

to want to know why.” However, the ultimate customer in front of the ATM or POS device does not necessarily want to know why the transaction failed. While declines are annoying, typically customers just take out another payment method, and do not ask their bank for an explanation.

CUT ON THE BIAS

The extent to which machines can be biased is an old issue in machine learning. Psychologists have found that humans often tend to prefer simpler models for explaining things in the world. Based on these simple models, they then generalise and draw conclusions for other situations that are new to them. Similarly, simpler statistical models often have a high bias, which enables them to generalise better.

The aim of a statistical model is not to most accurately predict outcomes for those in the sample. It is to achieve the best predictions for those outside the sample. By adding more and more predictor variables, the modeller is adding information. However, it may be the very information, enabling them to more accurately predict aspects about the existing sample, or even peculiar to the sample. These are not aspects which apply more generally to the population. It may seem counterintuitive that having less information can lead to better predictions. Yet simpler models avoid the problem of ‘overfitting’, where they follow the learning data too closely, and perform poorly when confronted with new situations.

“From a human language point-of-view, bias always has a notion of being prejudiced and

thus it is something you would want to avoid," says Dorndorf. "Algorithms per se I don't think can be biased. But algorithms are something that is invented by humans. Human decisions about data that you feed into algorithms and about the design of algorithms are, of course, subject to human bias."

The issues around causation and bias are a moot point. End customers may not be asking right now why their transaction has been declined. But how long will it be before they are asking why they have been refused service? Why are they being charged higher prices? And is this fair? One party keeping a watching brief on the use of algorithmic decisioning is the regulator. This is to ensure groups of customers are not being unfairly penalised. Also, that companies are handling personal data appropriately.

The UK Financial Conduct Authority stopped short of launching a full inquiry into the use of Big Data in the insurance industry in September 2016. But the watchdog warned insurers that they must comply with data protection regulations and privacy rules when using data from social media data or aggregator sites. Facebook also waded into the debate about the use of social media in the insurance industry last year. It blocked UK car insurer, Admiral, from using posts on its site to make decisions about pricing. Admiral had hoped that the language used in Facebook posts would provide clues about driving style.

Understanding and being able to explain the workings of machines, whether to improve decisioning or avoid the perception of unfairness, will be seminal to increasing the uptake of machine learning techniques both within the payments industry and beyond.

A SMARTER FUTURE?

Advances in artificial intelligence and machine learning have been driven in part by cheap access to computing power. Microsoft, Google, Facebook et al are good at artificial intelligence, partly because they have large swathes of computers to put behind the task. The machines will get smarter. The falling cost of computer power and data storage, increasing availability of data and advancements in AI techniques almost guarantee this. The production of models and calculation of results will become much better, quicker and cheaper.

Nick Bostrom, founding director of Oxford University's Future of Humanity Institute, believes that machines will replicate and surpass the intelligence that gives humans their dominance within our lifetime. Certainly within the lifetime of our children. Therefore AI will have a broad impact on humanity and, to Professor Hawking's point, we do not know exactly how. Technology is morally neutral, so we cannot guarantee that only the good guys will have AI or build it. As to the technology itself, how can we be sure that its intentions are good? Will there be an off button to ensure that AI does not circumvent human control? The technology giants, including Amazon, Facebook, Google, IBM and Microsoft, have set up a body to conduct research into ethical questions about how AI affects humanity. It is known by the slightly Orwellian name of the Partnership on Artificial Intelligence to Benefit People and Society, or PAIBPS for short.

Just as machines will get smarter, so will humans. They will understand the technology better and incorporate it more into day-to-day business processes. With regard to the payments industry in particular, as Hochmuth from INFORM maintains, there is a "huge opportunity to combine conventional decision-making strategies, human expertise and machine learning to improve processes." Machine learning will create a ripple effect as with all innovation. It will be a technology that enables a new raft of business cases and industries, including some we cannot possibly think of at present.

The future may be happening faster than we expect. Self-driving cars are a good example. Who would have thought that they would be on the road ten years ago? If machines can think, learn and communicate with other machines, how soon before they replace humans altogether? This dystopian view of the relationship between man and machines is not new. Nor is the discussion about machines in the workplace. Perhaps it is as US physicist Arno Penzias said. If you don't want to be replaced by a machine, don't try to act like one. As much as machine automation may make humans redundant, there are also opportunities. The future will involve new jobs that have yet to be created, as well as old

ones which cannot easily be replicated by machines. Cathy Davidson, co-director of the annual MacArthur Foundation Digital Media and Learning Competitions, claims that 65 percent of children will do jobs that do not exist today.

Machines will get smarter, so will humans. Yet the smart future will be powered by data. This really is the new oil on which the modern economy will run. This is not lost on European regulators. They are planning the biggest shake-up of data protection legislation in two decades. Leaving aside whether the regulation anticipates the future or is playing catch-up on present practices, the EU General Data Protection Regulation (GDPR) is already casting its shadow forward.

The regulation facilitates opportunities for data use and data sharing. However, it also penalises the mis-use or loss of data with fines of up to €20 million or four percent of global annual turnover. The regulation also includes provisions on the clear and affirmative consent to the processing of personal data, clearer language around privacy policies, the right to be forgotten, data portability and breach notification.

Attitudes to data at an individual and societal level will be fundamental to a smarter future. A cultural shift in the way we regard data, create trust and build business cases is already underway. The machines are rising, but the future will be determined by data and human attitudes.

