


Artificial intelligence is growing up fast—what's next for thinking machines?

February 6 2018



Artificial intelligence. Credit: The District

Our lives are already enhanced by AI – or at least an AI in its infancy – with technologies using algorithms that help them to learn from our behaviour. As AI grows up and starts to think, not just to learn, we ask how human-like do we want their intelligence to be and what impact will machines have on our jobs?

We are well on the way to a world in which many aspects of our daily

lives will depend on AI systems.

Within a decade, machines might diagnose patients with the learned expertise of not just one doctor but thousands. They might make judiciary recommendations based on vast datasets of legal decisions and complex regulations. And they will almost certainly know exactly what's around the corner in autonomous vehicles.

"Machine capabilities are growing," says Dr. Stephen Cave, Executive Director of the Leverhulme Centre for the Future of Intelligence (CFI). "Machines will perform the tasks that we don't want to: the mundane jobs, the dangerous jobs. And they'll do the tasks we aren't capable of – those involving too much data for a human to process, or where the machine is simply faster, better, cheaper."

Dr. Mateja Jamnik, AI expert at the Department of Computer Science and Technology, agrees: "Everything is going in the direction of augmenting human performance – helping humans, cooperating with humans, enabling humans to concentrate on the areas where humans are intrinsically better such as strategy, creativity and empathy."

Part of the attraction of AI requires that future technologies perform tasks autonomously, without humans needing to monitor activities every step of the way. In other words, machines of the future will need to think for themselves. But, although computers today outperform humans on many tasks, including learning from data and making decisions, they can still trip up on things that are really quite trivial for us.

Take, for instance, working out the formula for the area of a parallelogram. Humans might use a diagram to visualise how cutting off the corners and reassembling it as a rectangle simplifies the problem. Machines, however, may "use calculus or integrate a function. This works, but it's like using a sledgehammer to crack a nut," says Jamnik,

who was recently appointed Specialist Adviser to the House of Lords Select Committee on AI.

"When I was a child, I was fascinated by the beauty and elegance of mathematical solutions. I wondered how people came up with such intuitive answers. Today, I work with neuroscientists and experimental psychologists to investigate this human ability to reason and think flexibly, and to make computers do the same."

Jamnik believes that AI systems that can choose so-called heuristic approaches – employing practical, often visual, approaches to problem solving – in a similar way to humans will be an essential component of human-like computers. They will be needed, for instance, so that machines can explain their workings to humans – an important part of the transparency of decision-making that we will require of AI.

With funding from the Engineering and Physical Sciences Research Council and the Leverhulme Trust, she is building systems that have begun to reason like humans through diagrams. Her aim now is to enable them to move flexibly between different "modalities of reasoning", just as humans have the agility to switch between methods when problem solving.

Being able to model one aspect of human [intelligence](#) in computers raises the question of what other aspects would be useful. And in fact how 'human-like' would we want AI systems to be? This is what interests Professor José Hernandez-Orallo, from the Universitat Politècnica de València in Spain and Visiting Fellow at the CFI.

"We typically put humans as the ultimate goal of AI because we have an anthropocentric view of intelligence that places humans at the pinnacle of a monolith," says Hernandez-Orallo. "But [human intelligence](#) is just one of many kinds. Certain human skills, such as reasoning, will be

important in future systems. But perhaps we want to build systems that 'fill the gaps that humans cannot reach', whether it's AI that thinks in non-human ways or AI that doesn't think at all.

"I believe that future machines can be more powerful than humans not just because they are faster but because they can have cognitive functionalities that are inherently not human." This raises a difficulty, says Hernandez-Orallo: "How do we measure the intelligence of the systems that we build? Any definition of intelligence needs to be linked to a way of measuring it, otherwise it's like trying to define electricity without a way of showing it."

The intelligence tests we use today – such as psychometric tests or animal cognition tests – are not suitable for measuring intelligence of a new kind, he explains. Perhaps the most famous test for AI is that devised by 1950s Cambridge computer scientist Alan Turing. To pass the Turing Test, a computer must fool a human into believing it is human. "Turing never meant it as a test of the sort of AI that is becoming possible – apart from anything else, it's all or nothing and cannot be used to rank AI," says Hernandez-Orallo.

In his recently published book *The Measure of all Minds*, he argues for the development of "universal tests of intelligence" – those that measure the same skill or capability independently of the subject, whether it's a robot, a human or an octopus.

His work at the CFI as part of the 'Kinds of Intelligence' project, led by Dr. Marta Halina, is asking not only what these tests might look like but also how their measurement can be built into the development of AI. Hernandez-Orallo sees a very practical application of such tests: the future job market. "I can imagine a time when universal tests would provide a measure of what's needed to accomplish a job, whether it's by a human or a machine."

Cave is also interested in the impact of AI on future jobs, discussing this in a report on the ethics and governance of AI recently submitted to the House of Lords Select Committee on AI on behalf of researchers at Cambridge, Oxford, Imperial College and the University of California at Berkeley. "AI systems currently remain narrow in their range of abilities by comparison with a human. But the breadth of their capacities is increasing rapidly in ways that will pose new ethical and governance challenges – as well as create new opportunities," says Cave. "Many of these risks and benefits will be related to the impact these new capacities will have on the economy, and the labour market in particular."

Hernandez-Orallo adds: "Much has been written about the jobs that will be at risk in the future. This happens every time there is a major shift in the economy. But just as some machines will do tasks that humans currently carry out, other machines will help humans do what they currently cannot – providing enhanced cognitive assistance or replacing lost functions such as memory, hearing or sight."

Jamnik also sees opportunities in the age of intelligent machines: "As with any revolution, there is change. Yes some jobs will become obsolete. But history tells us that there will be jobs appearing. These will capitalise on inherently [human](#) qualities. Others will be jobs that we can't even conceive of – memory augmentation practitioners, data creators, data bias correctors, and so on. That's one reason I think this is perhaps the most exciting time in the history of humanity."

Provided by University of Cambridge

Citation: Artificial intelligence is growing up fast—what's next for thinking machines? (2018, February 6) retrieved 20 April 2024 from <https://techxplore.com/news/2018-02-artificial-intelligence-fastwhat-machines.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.