Ctrl Alt Del*iberate*

Artificial intelligence is changing the world in ways human intelligence has struggled to anticipate. But are we ready for the rise of the robots? **Anthony Sharpe** moves to the algorithm

In May last year, Anthony Levandowski – an engineer who built Google's self-driving car, co-founded an autonomous-truck company, and found himself at the centre of a massive lawsuit – filed papers to the United States IRS related to his founding of a church called Way of the Future. The intention of the church is: 'the realisation, acceptance, and worship of a godhead based on artificial intelligence developed through computer hardware and software'. According to the organisation's website, 'Way of the Future is about creating a peaceful and respectful transition of who is in charge of the planet from people to people + "machines".' ➔

Beyond demonstrating that Silicon Valley types should probably do a little less navel-gazing and a lot less cocaine, and moreover that you really can start a religion based on anything (Way of the Future was granted tax exemption three months later), it's a sign of the times: AI – whereby computer programmes use complex algorithms to perform tasks usually requiring human intelligence – is changing our lives in unforeseeable, exponentially more numerous ways.

If you haven't recognised that fact, it's likely because of something called the 'AI effect', a process whereby as any AI-based technology becomes sufficiently integrated into people's lives, they stop seeing it as intelligence rather than computation. But you've probably used AI several times already today. Your email spam filter? Powered by AI. Google Maps finding you the best route to your hotel? AI. The creepy way Facebook recognised faces in images you upload? AI. The hulking, emotionless man bearing an uncanny resemblance to the former governor of California who you drove past on your way to the airport? AI.

'I'M AFRAID I CAN'T LET YOU DO THAT, HAL'

Of course, while some folk may tend to take the increasing ubiquity of AI for granted, others are rather more concerned about its increasing power. The technological singularity refers to the idea that once artificial super-intelligence (ASI) is created, it will rapidly develop at an uncontrollable pace, with implications potentially beyond our imagination. The idea of ASI gone rogue has been a fiction staple for years, from *2001: A Space Odyssey* to *The Terminator* and *Avengers: Age of Ultron*. But this isn't just the stuff of science fiction anymore; some very smart people have voiced concerns about this, including Bill Gates, Elon Musk and the late Stephen Hawking.

One smart person who isn't too worried, however, is Geoff Nitschke, a senior lecturer in the Department of Computer Science at UCT performing research in a host of AI-related fields including evolutionary computation, artificial neural networks and neuro-evolution. He's currently working on evolutionary robotics, 'which combines AI and robotics, specifically taking biological (evolutionary) systems as inspiration for machine learning algorithms that are tested in real (and simulated) robots', and sees fears of a dystopic future as absurd.

'Many will argue that one of the reasons future super-intelligent machines pose a threat to us is that they, as ➔

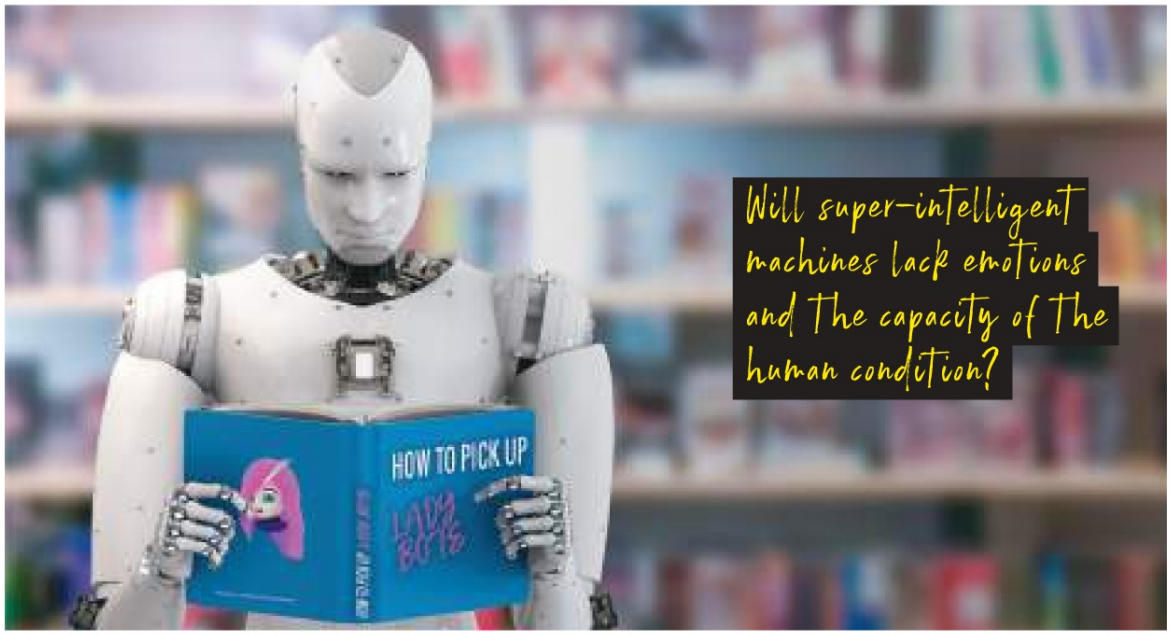

machines, will lack emotions and the capacity of the human condition,' says Nitschke. 'This is based on assumptions that we can isolate human intelligence from other human characteristics such as emotion, especially empathy.'

These considerations aside, there's the little snag that even contemporary supercomputers are not even remotely up to the task of handling this sort of software. 'The emergence of ASI is theoretically possible,' says Nitschke. 'However, our current computer technology, despite rapid pace advances, is insufficient to result in an ASI. It would take a fundamental change of computer architectures, languages, theories and computational methods. For example, major breakthroughs in neuromorphic computing, allowing brains of comparable complexity to that of humans to be created, could pave the way for a future with ASI.'

I HAVE NEURONS, THEREFORE I THINK, THEREFORE I AM

The thing is, human brains are complicated beasts, comprised of gazillions of neurons that exhibit a degree of plasticity – that is, they can grow, adapt and change over time. Compare this to the rigid constraints of silicon-based microchips, and it's clear there's little competition. That's where the field of neuromorphic computing comes in, the goal of which is to mimic the function of a biological brain using a new type of computer. 'This computer will be made up of artificial neurons and neural tissue engineered from replicated organic materials that make up biological central nervous systems,' explains Nitschke. 'In such a computer, the electrical impulses firing on organic connections between the artificial neurons would be the software, and the physical network of neurons would be the hardware.'

This, says Nitschke, might pave the way for the development of something akin to the intangible self-awareness of consciousness. 'However, if the computer remains a disembodied entity, like a brain in a jar, then consciousness seems unlikely since thousands of years of evolution of brains together with our bodies seems to have played a key role in developing the capacity to problem solve, reason, speak, experience a lifetime of events and perhaps to even be self-aware.'

This evolutionary factor is cited by others as a constraint on the possibility of a 'hard take-off' scenario: where an ➔

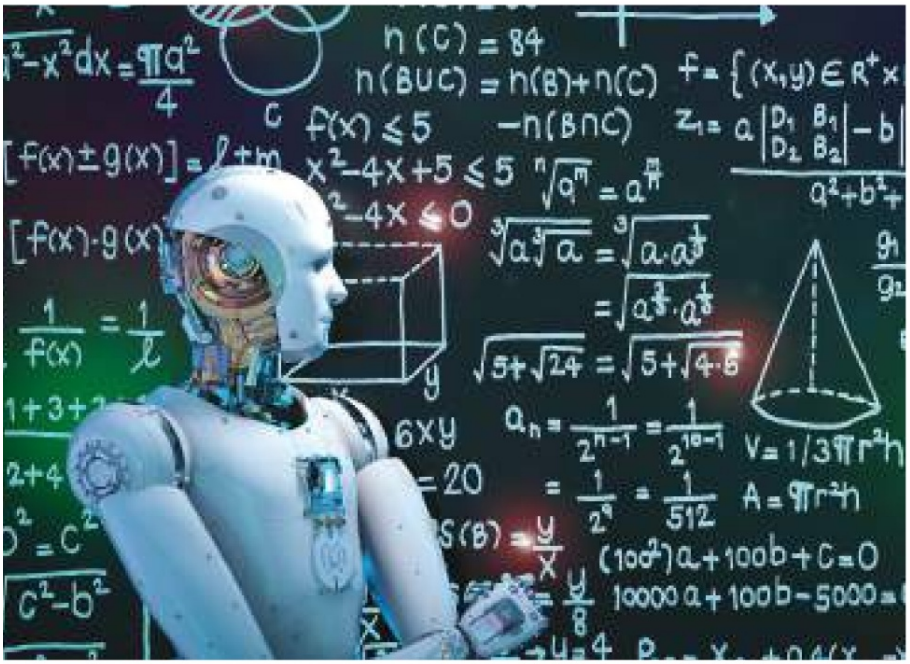

Communicating what AI is doing and can do is difficult

AI takes over the world in a matter of hours or days. Moreover, think about how often you upgrade your cellphone. Now imagine if you had to upgrade your *brain*.

THINKING INSIDE THE BOX

If artificial intelligence faces trust issues, it's probably largely because most people don't understand it (humans have a pretty well-documented history of fearing that which they don't understand). But here's the kicker: even programmers don't always understand why AIs arrive at the conclusions they do – they give inputs seeking certain outputs, but the inner workings of the system are opaque. This is known as the 'black box' concept in machine learning.

A major branch of research looking to address this is that of explainable AI (XAI), whereby the machine's actions can easily be understood by humans. The importance of this is underlined by the increasing use of AI in credit-scoring, profiling, medical and military applications.

Frans Cronje is managing director and co-founder of DataProphet, a Cape Town-based company that uses machine learning to optimise process parameter control in manufacturing. He sees XAI as a critical part of the company's message. 'AI historically and currently often loses trust by overpromising and underdelivering,' says Cronje. 'We need to be careful in communicating clearly what it is good at and what it is bad at. Right now, it is too easy to imagine AI to be human-like intelligence where, in fact, it is far away from that.'

He's quick to point out, however, that it's not an easy thing to do. 'Communicating what AI is doing and can do is difficult. One of the underlying principles of the AI theory is that through allowing AI to develop complex understandings of the task it is working on, it necessarily needs to be more complex. This means the better the AI is at a task, the harder it is to explain what it is doing. Communicating that relationship is difficult.'

Understanding AI is also a key part of regulating it, but those tasked with developing legal frameworks often don't understand the technologies they're dealing with. Nitschke believes AI experts should be involved every step of the way. 'Consider recent controversies about the use of algorithms to automate or assist with all kinds of jobs around the world, everything from ethnic bias in facial-recognition systems used by law enforcement, or internet bots used to write and disseminate fake news articles and influence political events,' he says. 'AI experts are indispensable for explaining the how and why, and correcting such AI applications so they are suitably used.'

WORK LIKE A MACHINE

Beyond the issues of understanding why it does what it does, what we should be more conscious of is how AI is going to affect our economies and the way we work. While history shows that technological advancements have ➔

traditionally usually resulted in net job gains, economists aren't so sure this time round.

A study published in March this year by the Organisation for Economic Co-operation and Development of 32 of its member states found that automation will impact almost half of all jobs, but that the risk isn't even.

Only 14% of jobs are highly automatable, while another 32% have a 50–70% chance of being significantly affected by automation, in a way that changes their skill requirements. Of concern is the report's finding that, despite widespread concerns that the AI revolution will do to white-collar jobs what the last industrial revolution did to blue-collar ones, lower-skilled jobs are the most at threat: 'Automation is found to mainly affect jobs in the manufacturing industry and agriculture, although a number of service sectors, such as postal and courier services, land transport and food services are also found to be highly automatable. The occupations with the highest estimated automatability typically only require a basic to low level of education. At the other end of the spectrum, the least automatable occupations almost all require professional training and/or tertiary education.'

It isn't going away; it's only going to get smarter. What people need to do now is adapt, prepare... and reap the rewards. AI is steadily working its way into almost every facet of our lives. Nitschke cites customer-service and behaviour-analysis bots, big-data analysis, computer games, biometrics for security and cybersecurity, and self-driving cars and drones as just five areas where AI is making a huge impact. On the local front, he also sees great potential for AI in the education system, such as with apps that provide a tailored learning experience to students. 'The AI role in such "teacher" applications would be that the application adapts to the competency and rate of learning of each learner, providing demonstrations, problems and solutions for an individualised learning experience,' says Nitschke. 'This means the AI application would be itself learning what each learner needs to do and at what pace.'

Is AI going to take over the world? It just might – but perhaps not in the ways we expect. 'The current state of research is still only at the point of building very narrow, and very specific, AIs that can perform just one task, and that cannot reason through environments they haven't seen before,' says Cronje. 'When an AI can have an argument with you, then we will need to worry.' ■

pictures: Phonlamai Photo/shutterstock.com, supplied

FOUR AFRICAN AI STARTUPS TO WATCH

- **Aerobotics:** This local firm uses drones (pictured right) and AI to scan farms, providing farmers with valuable data on crop health and feedback in order to maximise yields and reduce costs. *aerobotics.io*
- **Clevva:** Stellenbosch-based Clevva helps companies by programming the logic they use to make decisions into virtual advisor apps, or chatbots, which assist staff and customers. *clevva.com*
- **Xineoh:** Consumer-targeted marketing is becoming the new norm, and this Bloemfontein-based startup uses a combination of AI and machine learning to predict consumer behaviour, identify business opportunities and price goods accurately. *xineoh.com*
- **Ubenwa:** This Nigerian startup has developed an app that uses machine learning to analyse the amplitude and frequency of a baby's cry to monitor for signs of birth asphyxia, the third-highest cause of death in children under five years of age. *ubenwa.com*