

EDUCATION: FUTURE FRONTIERS

OCCASIONAL PAPER SERIES

The AI Revolution

TOBY WALSH | PROFESSOR OF ARTIFICIAL INTELLIGENCE



An essay commissioned by the NSW Department of Education

ABOUT THE AUTHOR

Toby Walsh is Scientia Professor of Artificial Intelligence at Data61, University of New South Wales.



EDUCATION: FUTURE FRONTIERS is an initiative of the NSW Department of Education exploring the implications of developments in AI and automation for education. As part of the Education: Future Frontiers Occasional Paper series, the Department has commissioned essays by distinguished authors to stimulate debate and discussion about AI, education and 21st century skill needs. The views expressed in these essays are solely those of the authors.

We are in the midst of a revolution in which Artificial Intelligence (AI) is helping to transform our political, social and economic systems. AI will impact not just the workplace, but many other areas of our society like politics and education. As with comparable events in the past like the Industrial Revolution, the road ahead may be bumpy in parts. This paper catalogues a number of the ethical challenges posed by AI. It ends with implications for the way our education system might help prepare society for this time of change.

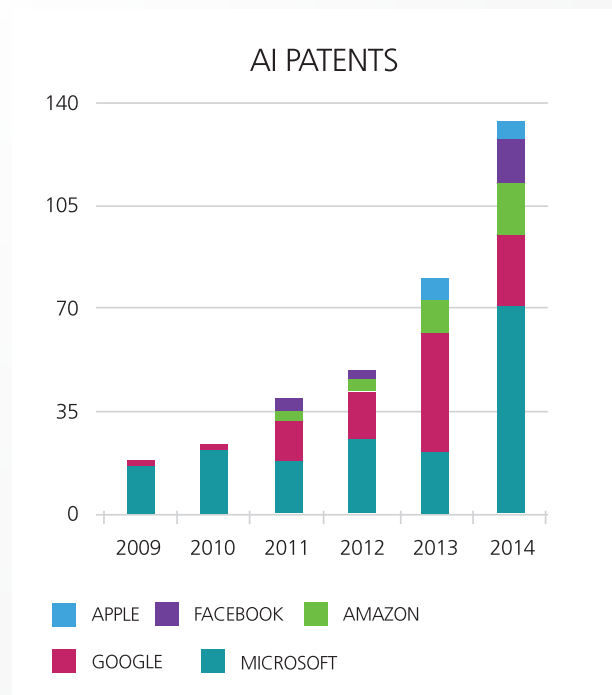
INTRODUCTION

Rapid progress is being made today in the field of AI and robotics. This is being driven by four exponential changes:

1. **Processing power:** Several decades of Moore's Law has doubled transistor counts every 18 months. Computational problems that were previously impractical are now becoming possible.
2. **Data:** The amount of data online is also doubling roughly every two years. Smartphones in particular, and the Internet of Things more generally, will continue this trend. This is providing data sets off which data hungry techniques like Machine Learning (ML) can work.
3. **Algorithms:** Many decades of research into algorithms is starting to pay off. AI methods like Deep Learning are leveraging improved processing power and larger data sets to deliver exponential improvements in performance.
4. **Funding:** Venture and other funds are pouring into the field. Over the last five years, the number of acquisitions of AI startups has increased 50 percent every year. The amount of venture funding being invested in AI startups is also doubling every two years.

Large companies like IBM and Toyota are investing billions of dollars into AI research. A number of countries like Canada and the UK have recently launched special government backed initiatives in AI. An arms race is taking place in Silicon Valley between the big technology companies. This can be seen, for instance, in their patent activity.

These four ingredients, exponential increases in computer power, data, algorithm performance and funding are fueling rapid advances in AI and robotics. Milestones are being passed in areas as diverse as transcription (computers now outperform humans at transcribing spoken Mandarin), diagnosis (computers outperform the best doctors at diagnosing pulmonary disease) and warfare (computers outperform the best human pilots in air to air combat).



These advances will likely transform the workplace. Many jobs will be automated. It will not just be blue-collar professions that are automated. Many white-collar jobs in areas like journalism, medicine and law are also under threat. As with any new technology, it is worth remembering that many new jobs will also be created alongside those that are destroyed. In addition, many jobs will be improved by automation, letting people focus on more creative, social and strategic aspects of the job whilst the machines do the routine and mundane. To understand the net effect, we must also take into account other factors like changes in demographics, the decreasing length of the working week, and the impact of globalisation.

I will not focus here on the challenges these changes to work pose to our education system. It will clearly require some significant changes in what we teach to equip students for these new jobs. The focus of this paper is on the other impacts this AI revolution will have on our economic, political and social systems, and on the many ethical challenges this will create. Given the speed of change, we need to start preparing soon.

WHERE WILL THIS ALL END?

We have no evidence to suggest machines will not eventually become smarter than humans¹. But building machines that are as smart or even smarter than us is unlikely to be an easy goal to achieve. It is a major scientific and engineering project. The human brain is one of the most complex systems we know. Trying to match it in silicon is not going to be easy.

Most experts in AI estimate it will take at least 50 years to get to human level intelligence in machines. Very few expect it will take much longer than a century. A serious research effort in "AI Safety" has begun recently to prepare for this moment and ensure that the goals of any such intelligent or super-intelligent machines align with those of humanity. Fears that the machines will take over anytime soon remain more the concern of Hollywood than the laboratory.

Before we get to machines as capable as humans, we will achieve what is called "weak AI", machines able to match or outperform humans in narrow tasks. Indeed, we have already done so in domains like playing chess or the ancient Chinese game of Go². Such weak AI already poses many ethical challenges. In fact, weak AI will often pose more challenges than super-intelligence. It will, for instance, result in systems that fail in unexpected ways. And, as has already been seen with the first fatal Tesla crash, it will likely lead to systems that humans trust too much.

AUSTRALIAN AI

Australia is one of the countries close to the front of this revolution. Australia punches above its weight in AI research. In August 2017, Australia hosts both the leading Machine Learning conference (ICML 2017) and the leading Artificial Intelligence conference (IJCAI 2017). A reflection of Australia's standing internationally is that Australia is the first country outside North America to have hosted the IJCAI conference for a second time. In addition, there is a healthy startup community in

¹Alan Turing refuted many of the common objections to intelligent machines in his seminal 1950 MIND paper which helped launch the field of artificial intelligence.

²In 1997, Gary Kasparov who was then reigning world champion at chess was beaten by IBM's Deep Blue computer. In 2016, Lee Sedol who is one of the world's best players at Go was beaten by Google's AlphaGo program.

AUSTRALIA HAS A NECESSITY TO BE AT THE FRONT OF THIS REVOLUTION. WE HAVE A HIGH WAGE ECONOMY, AND MANY LOW WAGE NEIGHBOURS.

Sydney, Melbourne, Brisbane and elsewhere fielding AI technologies. And there are several industrial labs in Australia like Data61 and IBM Research with an excellent track record of transitioning AI technologies into practice.

Australia has several natural advantages in this space. Our mining industry is already one of the most automated on the planet. Mines are an excellent place in which to develop robotics and automation, bringing both immense financial and safety benefits. Our finance sector is also well placed to take advantage of Artificial Intelligence. The ASX leads the world in the exploitation of new technologies like blockchain. Australia also has a number of other sectors like medicine, higher education and transport likely to be amongst the first to be impacted by AI.

Australia has a necessity to be at the front of this revolution. We have a high wage economy, and many low wage neighbours. We can only hope to

compete with the efficiencies brought about by greater automation. With commodity prices falling, automation has kept our mines competitive. Australia is also cursed by distances, both within the country and to other countries. Around 10 percent of our GDP goes into transportation costs. Autonomous vehicles could drastically reduce these transportation costs, and provide a means of reducing CO₂ emissions³. They can also help combat congestion that is choking our cities, save us from investment in expensive infrastructure, and provide personal mobility to disadvantaged groups like the elderly and the disabled.

The impact that AI will have on society will therefore likely be felt early on in Australia compared to many other developed countries. We will not have the luxury of observing what happens in the US or elsewhere. We will need to lead the way in adapting to the changes.

SOCIETAL CHALLENGES

I begin with several important challenges facing society that artificial intelligence raises: privacy, transparency, trust and fairness.

Privacy

Our privacy is increasingly under threat. As we shall see in many other areas, AI is both part of the problem, but also likely part of the cure. Both business and government can now use technology to get unparalleled insight into

³Autonomous vehicles will be able to drive more efficiently, but this won't lead to reduction in CO₂ emissions if we then drive more, live further from our work, consume more goods, etc.

our lives. With this comes great responsibility. It is much easier to end up with Big Brother if we have technologies, especially those based around AI, that can look into our lives at scale. The Admiral Insurance incident described here illustrates that companies are already experimenting with AI technologies that invade our privacy.

It is a little surprising that there has not been greater concern within society about the impact of technology on our privacy. The Snowden revelations should have been a wake-up call to society about the potential abuses. Few technologists were surprised that our emails were being read. Email is one of the easiest forms of communication that can be monitored. Unlike other forms of communication like the telephone or post, email is already in a form that is machine readable. In totalitarian states like East Germany, neighbour listened in on neighbour. But it is so much easier with AI technologies where computer can listen in on neighbour.

There are currently strong pressures on governments to invade their citizens' privacy. In the global war against terrorism, security agencies are struggling to find dangers hiding within society. It is tempting for them to use technologies like AI to look for potential threats. This raises many troubling ethical questions. If technology can make society safer, is it not worth the invasion of our privacy? Is our privacy invaded when only an algorithm and not a person looks at our data? If we have nothing to hide, should we care?

Transparency

Another area of concern is the transparency around decisions made about us as more and more of these decisions are handed over to machines. Many current AI technologies are black boxes, unable to explain how

Admiral Insurance

In November 2016, this FTS100 car insurance company announced a project to offer cheaper car insurance to young drivers. By reading people's Facebook pages using natural language processing (NLP) algorithms, they wanted to identify those new drivers most likely to be a good insurance risk. Following public outcry, Facebook shut the project down claiming it violated their terms of service.

Several lessons can be learnt from this incident. As is often the case, AI is both part of the problem and potentially also the cure. On the one hand, AI technologies - in this case NLP - enabled the invasion of people's privacy. On the other, AI technologies could also enable the individual to control precisely what government and business know about them. The incident highlights that technology creates new opportunities in advance of the development of suitable laws or norms. Should companies be able to "discriminate" on the price of your insurance based on your Facebook posts? Can companies be simply left to regulate themselves in this arena?

they come to particular decisions. For example, one of the most fashionable and successful AI technologies currently is Deep Learning. This has been used in tasks as diverse as detecting skin cancer, pricing insurance and predicting crime. But Deep Learning cannot provide a good explanation for its decisions. Deep Learning uses a complex network of "artificial" neurons, one triggering another. In addition, how this network is connected and behaves depends on the massive amount of data used to train the network. Describing the network, the triggering decisions and training data likely gives little insight into a particular decision.

Photos App

In July 2015, a news story broke that Google's app had automatically labelled a black couple as "gorillas". The app had previously labelled dogs as "horses". Google's error was not unique. Other tech companies have developed racially biased imaging software. Flickr tagged black people as "animals" and "apes". In Flickr's case, they also labelled white people as "apes". And HP's webcams were shown to be able to track white faces but not black ones.

Google quickly fixed the error, not by having the program correctly label gorillas, but by removing the "gorilla" label altogether. In this case, the issue was identified and fixed quickly. But there are many other areas where algorithms may be making similar mistakes without us realising. In areas like credit risk assessment, job matching, online dating and product recommendation, algorithms are making decisions which impact our lives with very little transparency about how they work or why they make particular decisions.

As the image labelling examples above illustrate, we can unintentionally end up with damaging biases. Without transparency, we may never realise that certain groups are being discriminated against. In Europe, awareness about this issue is perhaps more advanced than elsewhere. In May 2018, the General Data Protection Regulation comes into law. This requires that personal data be processed transparently, that meaningful information be provided about the logic involved in any automated decision making, and that individuals have the right not to have decisions about them made entirely automatically. Such a law may become necessary here too.

There are also areas like national security where transparency is undesirable. We do not want terrorists to be able to know how threats are identified and monitored. A new scientific field at the intersection of game theory and computer science called "security games" is under development to enable computers to allocate limited security resources in an optimal way that is unpredictable.

MANY CURRENT AI
TECHNOLOGIES ARE
BLACK BOXES, UNABLE
TO EXPLAIN HOW THEY
COME TO PARTICULAR
DECISIONS.

COMPAS

In May 2016, the non-profit investigative news agency ProPublica revealed that the COMPAS program, used by judges in 20 of 52 states in the US to help decide parole and other sentencing conditions, was racially biased. COMPAS uses machine learning and historical data to predict the probability that a violent criminal will reoffend. Unfortunately it incorrectly predicts black people are more likely to re-offend than they do. And it incorrectly predicts that white people are less likely to re-offend than they do.

With work, we could improve the program to predict correctly whether someone is likely to re-offend. But how do we know when we can trust such a program? And there remains the deep philosophical question of whether machines should decide on who is locked up. Are there some decisions we should perhaps not hand over to machines, even if they make them better than us?

Trust

Closely connected to concerns about transparency are concerns around trust. How do we know when to trust a machine? What information provided by machines can we trust? Will we perhaps trust machines too much? AI will likely make these issues more problematic. When we observe a computer performing intelligently on one problem, we often tend to suppose it will work equally well on another. In reality, however, AI remains very brittle. Our smart computers can be surprisingly dumb when the problem changes even slightly.

In safety and security critical areas, there are already well developed tools and techniques for verification and validation of computer systems. Unfortunately, these tools

TAY chatbot

In March 2016, Microsoft released the TAY chatbot onto the internet. TAY was designed to learn from the tweets coming from its teenage audience and to speak therefore like a teenage girl. Less than 24 hours later, Microsoft were forced to disconnect TAY as she had been taught to be racist, sexist and highly offensive.

In putting TAY onto the internet, Microsoft made a number of fundamental mistakes. They should have put a profanity filter on the input and output of TAY. And, they should not have left TAY to learn from the twitter-sphere without any checks. If a technology company like Microsoft makes such mistakes, you can be sure that we will see lots of similar mistakes from other companies in the near future.

TAY highlights a number of ethical challenges. Do chatbots have freedom of speech? Who is responsible for the actions of an AI program, especially when it uses Machine Learning and so is a product of both its initial code and the training data? How do we guarantee the behaviour of programs involving Machine Learning?

and techniques struggle to scale to complex AI systems, especially those that learn and change, and that interact with a complex environment. We are even challenged in defining what properties machines should have for us to trust them. What, for example, does it mean that an algorithm is racially unbiased?

Despite what high-tech companies like Google might have us believe, algorithms especially those using Machine Learning, can be biased. Algorithmic discrimination will start to trouble society increasingly. If we are not careful, many of our hard fought rights against racial, religious, sexual, age and other types of discrimination will be lost to machines that are not transparent, and that we should not trust.

Fairness

With economical, environmental, and societal pressures mounting, countries are struggling to use their limited resources more fairly. As we start to hand decisions over to AI systems, we will want to ensure that they act fairly. In fact, computation can actually improve what they do. We can, for instance, have the system compute outcomes which are both fair and efficient.

Building AI systems that act fairly raises a number of ethical questions. What does fairness formally mean? For example, suppose we write a program to allocate organs to patients. How do we fairly treat patients of different blood type and age? At the same time, how do we fairly treat the different hospitals and states? How do we treat different ethnic groups fairly, recognising that some might be disproportionately present on the waiting list? And can we be fair to all these different actors simultaneously?

POLITICAL CHALLENGES

Other aspects of our society will be affected by AI. We are already witnessing the impact of algorithms on politics and political debate. Cambridge Analytica, the data driven political marketing company behind both the Trump Presidential campaign and the Pro-Brexit vote, is looking to expand into Australia. Using psychological data derived from millions of Facebook users, Cambridge Analytica tries to identify key swing voters. When do we cross the line from convincing to manipulating? Is a technological arms race between parties to target voters destructive to democracy? If we use algorithms to influence voters at manipulating scale, does it threaten our very democracy?

Another area of concern is fake news. Following Trump's election, many commentators suggested that fake news might have had a significant impact on the result. Facebook initially denied responsibility for the propagation of fake news. However, in February 2017, Facebook CEO and co-founder Mark Zuckerberg accepted some

Facebook

In June 2014, news broke that Facebook had secretly run an A/B experiment, not to improve their product, but to see if they could change the mood of their users. They altered the number of positive and negative posts in the news feeds of 689,003 randomly selected users. Users with more positive posts were observed to post more positively than users shown more negative posts. No ethics approval was sought for the experiment.

Not surprisingly, Facebook apologised. Several fundamental issues remain. When running tests involving the public, should companies like Facebook and Tesla have to face the same ethical hurdles that researchers have to face at universities? Should companies be allowed to manipulate people's emotions like this? Do we need more regulation of technology companies? Is government giving them too free a hand?

responsibility in an open letter. Interestingly, many of the suggestions he proposed for tackling fake news involved using AI. This is not too surprising. The only way you could filter hundreds of millions of posts each day is with AI-based natural language processing technologies.

A third political concern is freedom of speech. Who or what is responsible for the messages that machines produce? This is especially difficult to decide when Machine Learning is involved. The program may produce output that is very unexpected. What if the machine incites racism? How free is human speech when it is drowned in a sea of machine voices? It is estimated that over three quarters of Trump's twitter traffic during the last Presidential election were fake supporters, Twitter bots that artificially boosted the Trump message.

HUMANITARIAN CHALLENGES

It ends with a major humanitarian and ethical challenge introduced by AI. There is an arms race underway today to develop lethal autonomous weapons, or as the media like to call them, “killer robots”. This will be the third revolution in warfare, after the invention of gunpowder and nuclear weapons. There are many reasons to fear this change. It will herald a step change in the speed and efficiency with which we can kill the other side. It will destabilise the current geopolitical order. These will be weapons of terror, and of mass destruction. Unexpected feedback between swarms of such systems may trigger unwanted wars just as we see “flash crashes” in the financial markets triggered by interactions between trading algorithms. As a result, many AI researchers and NGOs like Human Rights Watch are now campaigning for a pre-emptive UN ban on such weapons.

Lethal autonomous weapons raise a whole host of ethical challenges. How do we build robots that behave ethically? Could robots be built to follow international humanitarian law (IHL)? Could they distinguish adequately between combatant and civilian in the fog of war as required by IHL? Who is responsible for their actions? How do we prevent them being hacked to behave unethically? Should machines be given the right to make life or death decisions? Should there also be a human “in the loop”? Many of these ethical decisions will be faced when we let robots into other parts of our lives. It is just that the setting of the battlefield makes the ethical choices even more stark.

HISTORICAL LESSONS

This is not the first technological revolution that has affected society so we might look for lessons that can be learnt from history. Perhaps the closest parallel is the Industrial Revolution. This liberated us from the limitations of our muscles, transforming the nature of work. Before the Industrial Revolution, much of the world’s population was occupied in farming. Automation replaced many of these jobs so that today just a few percent of the workforce is left in agriculture. New jobs were, however, created in factories and offices that employ those displaced from the fields.

In the Industrial Revolution, we still had a cognitive advantage over machines. It is less clear what advantages we will maintain over the machines this time. There is another reason that this time is different. Not because this time is special, but rather because last time was very special. At the time of the Industrial Revolution, the world took several large shocks which helped society to adapt to the change. Two World Wars and the intervening Great Depression set the stage for what economists are now starting to recognise as an unusual reversal in inequality.

The introduction of the welfare state, of labour laws and unions, and of universal education began a period of immense social change. We started to educate more of the workforce, giving them jobs rather than allowing machines simply to make them unemployed. At the same time, we provided a safety net for many, giving them economic security rather than the workhouse when machines made them unemployed.

WE WILL NEED AN INFORMED POPULATION TO NAVIGATE THIS FUTURE ... A CITIZENSHIP EDUCATED IN ETHICS, SOCIETY AND CIVICS IS THEREFORE ESSENTIAL.

We might expect equally large societal changes will occur and will be needed for the coming AI revolution. A worrying lesson from history is that there was around half a century of pain at the start of the Industrial Revolution during which prosperity for many in society went backwards. It took some time before society adapted so that technological progress improved the lives of many.

IMPLICATIONS FOR GOVERNMENT

Motivated by these ethical concerns and historical lessons, I will identify a number of implications for government. All concern education in one way or the other. This is because education is one of the most important and powerful tools at our disposal in adapting to the coming changes.

Teaching ethics, society & civics

In fifty years time, we may look back at the next decades as a golden age for ethics. In handing over many of our decisions to machines, we will need to make explicit in computer code many of our society's ethical choices. This will require us to have much greater clarity and consensus about what these ethical choices are.

With society under a period of significant change, we will also need an informed population to navigate this future, and to demand appropriate checks and safeguards. A citizenship educated in ethics, society and civics is therefore essential. The education system needs to prepare us for this future of "computational ethics".

Teaching creativity

One of the advantages that humans have over machines is our creativity. Computers struggle to be creative. Machines are excellent at doing the routine and repetitive, and poor at coping with change and unpredictability. In time, I expect that machines will become as creative and adaptable as humans. However, for the next few decades at least, we will have a significant edge over machines in this area.

A creative population will be able to keep itself employed and ahead of the machines. Even if machines can be creative, they cannot speak to the human experience: about love, death, and all the things that make us unique. A creative population will also be able to take advantage of the free time that

automation may give us. It follows that creativity can and should be taught more actively. If machines take over the sweat, this could leave us with the time to create the next Renaissance.

Developing emotional intelligence

Another advantage that humans have over machines is our emotional intelligence. Computers struggle to understand our emotions. And they have no emotional lives of their own. As with creativity, we are likely to have the edge over machines in jobs that require emotional intelligence for a long time to come. In addition, there will be an increasing value placed on social contact between humans. Emotional intelligence will therefore be increasingly important.

At present, our current education system focuses on lifting cognitive abilities. However, in some countries like Germany, attention is also given to improving emotional intelligence. Classes in Germany will often have both a teacher, focused on the children's cognitive development, and an educator, focused on their emotional development. This would be a good idea here too in Australia.

Universal lifelong learning

For many, education stops when they leave school or university. This is undesirable if we are to keep ahead of the machines.

We need to re-invent ourselves constantly, learning new technologies, and adapting to the unexpected changes occurring within society. This requires an education

system that gives us not just knowledge but learning skills, so we can learn throughout our working lives. We need to learn how to learn so that we can continue to learn even when we are no longer in a formal education environment like a school or university.

Government will need to support such lifelong learning, providing financial and other incentives to individuals and businesses to encourage re-skilling of the workforce. Ultimately, just as the Industrial Revolution made it essential that universal education was provided to the young, the AI Revolution will make it essential that education is provided to people at every age of their lives.

Sea of dudes

In Australia and the US, a major problem within the field of Computer Science in general, and especially within Artificial Intelligence, is the under representation of women. This has been nicknamed the "sea of dudes" problem⁴. The imbalance starts in secondary school. By the time university starts, it has become sufficiently extreme that any corrective measures merely put sticky plaster on the problem.

The under-representation of women in AI and robotics is undesirable for many reasons. Women will, for instance, be disadvantaged in an increasingly technically focused job market. It may also result in the construction of AI systems that fail to address issues relevant to half the population, and even to systems that perpetuate sexism. More initiatives are therefore needed to get young girls interested in STEM in general, and AI and robotics in particular. It will also be worth exploring why women

⁴This phrase was coined in 2016 by Margaret Mitchell, then an AI researcher at Microsoft Research and now at Google. Her phrase highlights the fact that only around 10% of AI researchers are women. Actually, she might have more accurately described it as "a sea of white dudes". Not only are most AI researchers male, they are also mostly white.

are better represented in other countries. For example, women make up 30% of undergraduates in engineering courses in Spain compared to just 19% in the US.

One robot per child

In the 1980s, the UK government kick-started computer literacy by introducing the BBC Model B computer into every school in the country. Many students also started to have access to low cost computers like the Sinclair ZX80. At the time, there was significant scepticism of the value in giving children access to personal computers. What could they possibly learn from having access to word processors, spreadsheets and computer games? Two decades later, the UK found itself at the centre of the billion dollar computer game industry. This is not a coincidence.

Providing one robot per child will likely have similar unexpected but valuable side-effects. It will, of course, have the primary effect of promoting literacy in AI and robotics. But it is hard to predict the secondary effects it will have. Perhaps Australia will become the centre of the industry which personalises robots? Or a major force in the robot entertainment business? It may even position Australia as a leading player in a new personal robotics industry that rivals the personal computer industry.

Any robots put into schools should have both software and hardware that is open so students can be creative with them. They should also come with tools to help students explore less technical issues like ethics and social relationships. There is evidence that access to robots, especially at an early age, can help bring girls into STEM.

Computational thinking

We need citizens in our society to understand the fundamental principles of computation. If we don't, a large section of the population will be greatly disadvantaged as much technology will simply be magic to them.

This doesn't mean we need to teach everyone to hack code. But we do want people to understand the building blocks of computation, to appreciate what can (and can't) be done, to abstract problems so that they can be automated, to decompose problem solving into a series of algorithmic steps, and to generalise to work across problem domains. These problem solving skills will become essential in many new jobs. Robots will offer an excellent platform on which to teach such computational thinking.

Open educational data

Data in government should be opened up so that outside parties can innovate. Education should be at the centre of this open data revolution.

It will take some political courage to put education data at the centre of an open government as this will, for instance, expose where the system is failing students. But there will be many benefits.

Education can become more evidence based. Parents and students can be more informed in their choices. Teachers can share best practice. Heads can identify areas in their schools needing improvement. Universities can target disadvantaged students who might not otherwise

benefit from higher education. And high tech companies like Google and IBM, as well as startups, can produce software optimised to actual learning experiences.

Government-wide thinking

My final recommendation is for a government wide report on how to prepare for the changes that AI and Robotics will bring to society.

These are technologies that will touch almost every aspect of our lives. They will require changes to the welfare state, our taxation and pension system, schools and universities, our legal system, police force and armed forces, our health care system, transportation and housing, even perhaps our political system. This is not a transformation where we can or should consider the different parts of government separately.

At the end of 2016, the White House Office of Science and Technology, and the Joint Committee on Science and Technology of the House of Commons and of Lords both published reports on the challenges posed by AI and robotics. The US report especially contains some valuable recommendations. However, neither addresses features specific to Australia like our particular demographics, our geographical isolation, or our urban characteristics.

The NSW Chief Scientist, Mary O’Kane was previously an AI researcher. She would therefore be an excellent person to chair such a report. The UK report recommended setting up a standing committee to monitor this area. Such a committee might be useful in Australia. Both reports also recommended more government investment

in the area. If Australia is to compete in the worldwide AI arms race, it is likely that both government and business in Australia will also need to invest more.

CONCLUSIONS

The AI Revolution will transform our political, social and economic systems. It will impact not just the workplace, but many other areas of our society like politics and education.

We need therefore to start preparing for this future. There are many ethical challenges ahead, ensuring that machines are fair, transparent, trustworthy, protective of our privacy and respect many other fundamental rights. Education is likely to be one of the main tools available to prepare for this future. A successful society will be one that embraces the opportunity that these technologies promise, but at the same time prepares and helps its citizens through this time of immense change. ➔

REFERENCES

Alan M. Turing. Computing Machinery and Intelligence, *MIND*, 59 (263): 433-460, 1950.

Preparing for the Future of Artificial Intelligence.

Executive Office of the President. National Science and Technology Council Committee on Technology. October 2016. https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf

Robotics and Artificial Intelligence. Fifth report of Session 2016-2017.

House of Commons Science and Technology Committee. September 2016. <https://www.publications.parliament.uk/pa/cm201617/cmselect/cmsctech/145/145.pdf>

Jeanette Wing. Computational Thinking. *Communications of the ACM*. 49 (3): 33. 2006.



© June 2017.