No Fate But What We Make:

Policy Directions for AI in New Zealand

January 2019

TOWARDS OUR INTELLIGENT FUTURE TE ARA MŌ TĀTOU ATAMAI O ĀPŌPŌ

aiforum.org.nz

AI forum NEW ZEALAND

Te Kāhui Atamai Iahiko o Aotearoa

sapere®

No Fate But What We Make:
Policy Directions for AI in New Zealand

Hayden Glass

January 2019

KEY POINTS

- AI technologies will have very big impacts on New Zealand's economy and society. Credible estimates say that the changes brought by AI will be as significant as the ongoing changes that have brought computers and internet into all aspects of business, government and everyday life over the last thirty years.
- Generally New Zealand will be a taker not a maker of AI technologies. The scale of the impacts, both positive and negative, and the speed with which they appear will depend on how quickly and deeply organisations of all kinds absorb and adopt AI technologies. There is a wide range of possible timeframes for the transformation. It might take forty years to reach widespread adoption.
- New Zealand has some strong pre-cursors for the take-up of digital technologies, and some recent successful history with other digital infrastructure projects. The population is well-educated, the labour market is flexible, and the New Zealand economy is less exposed to the threat of automation than other places.
- AI is a general-purpose technology, which means it can be applicable across many industries and many different uses. The opportunities for effective application are still being developed and brought to market, and it is not clear at this point which opportunities New Zealand should look to prioritise.
- In general, policy choices can have big impacts on the direction of technology change. Countries that act with speed and judgement are expected to generate greater benefits from AI. As with the technology itself, the relevant policy options around AI are still forming, and direct line of sight from investment in AI to moving the dial on tangible economic or social measures is yet to be empirically established.
- As at November 2018, eighteen countries have published national AI strategies. Popular policy options involve encouraging private sector adoption and use of AI, funding research and supporting commercialisation of new ideas, training and attracting people with the relevant skills, and exploring AI ethics.
- The New Zealand Government has yet to formally begin work. One question is how quickly New Zealand wants to move on adoption of AI into the economy and society relative to other places. We present some high-level options for responses in this paper based party on the strategies of other countries.
- There is also scope to encourage the use of AI to resolve broader societal challenges, and to lead the deployment of AI to generate wider social benefits, rather than being focused mostly on reducing the costs of human labour for organisations.
- Given the wide impacts, an ongoing societal conversation about the social licence for AI and the tradeoffs involved the wider use of technology will be needed. Sometimes specific applications of AI will give rise to societal concerns that will be reflected in regulatory control.

SUMMARY

The impacts of Artificial Intelligence (AI) technologies on New Zealand society are expected to be similar in scale and scope to those that have accompanied the ongoing adoption of computers and then Internet services into business, government and everyday life over the last thirty years. As a set of general purpose technologies, AI has an enormous number of possible applications. What is visible now is just the early stages of commercialising technology that has been in existence for many years. Only now is the data and computer processing power starting to become available to take advantage of them.

As these technologies continue to be developed and deployed there will be significant changes in how work is done, what jobs are available, and the division of labour between humans and machines. The powerful data processing and cheap prediction enabled by AI also offer the promise of new approaches to pressing social challenges, and to the conception and delivery of government services.

The impacts of AI in New Zealand depend on how organisations integrate AI into their processes, products and services. Based on previous technology waves, it could take in excess of forty years for the implementation of AI technologies to reach a mature state and it will take a long time for the full economic impacts to emerge. There is still an active debate among economists about the productivity impacts of the use of computers and the Internet, even as they have become ubiquitous.

Organisations in New Zealand will largely be users ("takers") rather than makers of AI technologies, just as for other types of modern technology, like cars, computers or televisions. The extent to which organisations outside the tech sector adopt AI will therefore determine the scale of the impacts of AI here and how quickly those impacts emerge. A small proportion of organisations will be leaders in the use of AI, and most firms will be slower to take advantage of the opportunities.

The exact future capabilities of AI are uncertain, but clearly software can already do a lot more human work than it is doing right now, and credible estimates are that even existing technologies can significantly increase organisational capabilities and business profitability in many ways. These benefits mean that standard commercial and organisational incentives will be the strongest forces driving adoption. As the capabilities of AI improve over time, the benefits available from effective use of AI will continue to grow.

In the context of likely rapid adoption, there is a set of choices to be made about the direction that these technologies should take and what problems to apply them to. New Zealanders are generally positive about the benefits of new technologies. But AI presents some new ethical and practical issues to grapple with, and if New Zealand wants to, for example, encourage the use of AI to solve big societal challenges rather than just doing many of us out of a job, then it would be wise to explore options to do that sooner rather than later.

More generally, there will need to be an ongoing conversation about the social licence for AI. And doubtless from time to time specific applications will be stumbled across where societal concerns about the impacts or uncertainties of a new technology mean its deployment is foreclosed, postponed or limited for organisations in this country.

New Zealand is typically in the top 20 but not the top 10 in global indices of innovation and competitiveness. New Zealand's large number of small firms, primary sector and tourism focus, and small domestic markets reduce the incentives and capability to innovate. But New Zealand also has some strong pre-cursors for the take-up of digital technologies like AI, and some recent successful history with other digital infrastructure projects in the form of the Rural Broadband Initiative (RBI) and Ultra-Fast Broadband (UFB).

Some argue that countries that are not in the vanguard of the use of AI technologies will experience sustained economic underperformance in the future relative to the global leaders.

Unsurprisingly, many countries are already focused on taking advantage of the economic and social opportunities. As at November last year, 18 countries had released some form of national strategy.

The New Zealand government has said that it is at work on a plan of action, and it has some existing support for some work on AI ethics. No doubt the plan will build from an informed understanding of New Zealand's assets

and industry structure and the actual challenges that AI can help us with. Exploring how to get some focus on the issues that will not be readily addressed by commercial operators seems like it would be helpful. And a systematic assessment of existing use of AI and how much public science money is already being allocated to relevant projects would also be a super useful contribution. Early evidence is that understanding of AI in the private sector or in government is modest and actual use of the technologies is limited at this point in New Zealand.

There is a whole host of possible actions available if faster adoption or government leadership was thought desirable. These include amplifying stories of the practical uses of these technologies, educating and attracting talented people, expanding the funding and clarifying the direction of research, lifting the prevalence of AI startups and venture funding, improving the availability of data, and coordinating academic, government and commercial actors in the development and use of AI in their products, services and daily activities. There is no shortage of prediction problems in government agencies, for example, that might benefit from machine learning approaches. Courtesy of sustained effort over many years, New Zealand has some very good basic data infrastructure for trusted government use.

There is some talk in the extensive commentary around AI of an "arms race" or a "war" between the major AI global players US and China. In practice, the benefits available from smart use of AI are more open to sharing than that terminology suggests. A useful next step would be for New Zealand to start to think about its niches in AI and to consider some policy actions that will influence the impacts of AI on the economy and society as these technologies continue to be developed and put into productive use.

INTRODUCTION

We were asked by the AI Forum of New Zealand to provide a short paper as an early part of its research project into the economic and social impacts of AI, and what national responses to the opportunities presented by AI could be, building on a previous project from 2018.

This paper summarises the results of a focused look at:

- What adoption of AI might look like – We review frameworks for thinking about technology adoption and acceptance, as well as some historic examples that can illuminate how quickly the transition might happen, where it might end up and the factors that matter.
- The expected economic and social impacts of AI adoption, including descriptions of potential future scenarios – This builds on the material we contributed to previous AI Forum work, and includes views on potential productivity and labour market impacts, and social applications of new technology.
- Broad directions for AI policy responses – We outline some high-level options for thinking about how New Zealand could respond to the opportunities and challenges of an AI-powered world, based in part on responses in other countries.

Of these three points, we say more in this paper about adoption, the social licence, and policy responses, because these seem to be areas that will need more engagement in the coming years. We say less about the impacts of AI because that picture is already well covered and there is not as much to add to previous work.

An exact definition of AI is not crucial to this paper. We can follow Parker (2018), who says it "refers to machines doing things that human beings would normally do". Previous AI Forum research (AI Forum 2018, page 26) says it "includes technologies that enable machines to learn and adapt, to sense and interact, to reason, predict and plan, to optimise procedures and parameters, to operate autonomously, to be creative, and to extract knowledge from large amounts of data". Or more technically it can be helpful to think about AI as a family of technologies. McKinsey Global Institute (2018 page 4) studies five broad categories: computer vision, natural language processing, virtual assistants, robotic process automation, and advanced machine learning.

AI is a subset of information technology in general. It is not always easy to separate different types of technology from each other. So, questions about the adoption, impacts and sensible policy responses for AI shade into those same questions for technology in general.

ADOPTION

Adoption is a measure of the extent to which businesses, government entities and non-profits (together: “organisations”) use AI solutions in their operations. As with technology in general, most organisations will be users, with organisations buying AI-enabled services provided by others and larger organisations also deploying AI technologies within their environments themselves. A subset of organisations will be makers, involved in developing AI solutions. New Zealand’s tech sector accounts for about five per cent of total employment, which gives a sense of the possible scale of AI users relative to makers. More broadly, New Zealand organisations will generally be in the position of using AI technologies developed elsewhere, much as it is with other sorts of technology today. There are substantial benefits available from effective technology use.

Theoretical models say that perceived ease of use and perceived usefulness are major influences on the adoption of technology (Davis 1989). Ease of use and usefulness are mediated by a range of other issues, including what is popular, and how appealing a technology looks, and objective issues like how relevant it is to one’s work, and what difference it actually makes to output (Venkatesh and Davis 2000). Others add that cost, social influences, and structural imperatives are also important, eg, within an organisation, most individuals have fundamental technology choices made for them (Lunceford 2009).

Building on these models, one simple way to think about adoption of technology is as a combination of what is technically possible for AI to do, what is socially or legally permissible and what is profitable (or in the case of non-profit organisations, what is effective in advancing towards their goals). This framework gives us a way to think in general about technology that is not yet widely available. The more useful, more welcome and more profitable it is, the faster one can expect adoption to be. But even in the case of technology that has all these characteristics, mass adoption will still take time.

FIGURE 1

Components of technology adoption

In what follows, we consider each of these three matters in turn. And then turn to the question of how long widespread adoption of AI might take, and to the evidence we have on the level of adoption so far.

What is possible

In this case, “possible” refers to tasks that AI can complete that were previously the preserve of humans. Practical recent examples of things that computers powered by AI can now do well that might previously have been thought impossible include winning at Go (a very complex board game),[1] tagging photos based on their content,[2] providing customer service through online chat,[3] painting,[4] creating human faces,[5] classifying cancers or detecting eye disease from imagery,[6 7 8] and translating news from Chinese to English.[9]

1 https://www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611/

2 https://techcrunch.com/2018/12/18/box-releases-skills-which-lets-developers-apply-ai-and-machine-learning-to-box-content/

3 https://www.forbes.com/sites/forbestechcouncil/2018/05/23/how-the-development-of-ai-has-advanced-the-technology-available-for-chatbots/#2823c993c213

4 http://time.com/5435683/artificial-intelligence-painting-christies/

5 https://motherboard.vice.com/en_us/article/mby4q8/these-people-were-created-by-nvidia-ai

6 https://www.nature.com/articles/nature21056

A summary by an expert committee in the United States says (EOP 2016b, page 7):

> *"Narrow AI underpins many commercial services such as trip planning, shopper recommendation systems, and ad targeting, and is finding important applications in medical diagnosis, education, and scientific research. These have all had significant societal benefits and have contributed to the economic vitality of the [United States]"*

There is widespread commentary that computers and robots using AI technologies are now or will soon be able to technically do a very large share of work that is presently done by humans.

- Hanson (2001) presents a theoretical economic model where organisations start by using computers for the tasks they are better than humans at, and as computer capability expands and investment continues, eventually organisations replace all human labour with machines.
- Empirical studies have been based on expert judgement on the proportion of jobs or tasks that computers or robots are predicted to be able to do. The two classic references are Frey and Osborne (2013) who predict that about half of all human jobs in the United States could be replaced by automation in the coming decades, and Artnz et al (2016) who say that about 9 per cent of jobs could be replaced. A higher proportion of jobs will change than be replaced because many tasks within jobs will be done more efficiently by computers. New Zealand workers are thought to be less at risk because a higher proportion of their jobs are in less automatable industries (Nedelkoska and Quintini 2018).
- Elliot (2017) says that 62 per cent of workers in OECD countries use the basic literacy, numeracy and problem solving skills tested in PIAAC every day at work with a proficiency that computers can or will soon be able to reproduce, and that only 13 per cent of workers use these skills on a daily basis with a higher proficiency. (The other 25 per cent of workers do not use these skills on a daily basis). PIAAC is a standardised global skills test.
- Shoham (2018) presents a range of data on AI performance on various tasks all of which show improved AI performance over time. The metrics include object detection in images (now superior to human performance), and time to train an AI programme to classify images (down from an hour in June 2017 to four minutes by November 2018), translation between languages, and determination of the grammatical structure of sentences. The authors summarise (page 6) "AI activity is increasing nearly everywhere and technological performance is improving across the board".

AI certainly offers the promise of completing routine tasks faster better or with more consistency than humans. AI will follow the exact programming, it does not tire, it can find connections in data that humans cannot, it is cheap to scale, it needs only electricity, Internet connectivity and data, and intelligence can be built into a wide range of physical objects.

Lee (2018) sees four types of AI from a consumer's perspective: "Internet AI" that already powers search and recommendation engines, "business AI" that improves performance on analytical tasks within organisations, "perception AI" which brings sound and visual inputs to computers, and "autonomous AI" like self-driving cars.

There is much focus on what Lee calls "business AI", in particular machine learning for data analysis, computer vision and natural language processing. These are AI capabilities that are reasonably well-developed and widely applicable to typical tasks that involve classification or prediction in organisations (McKinsey Global Institute 2018b).

But there is quite a contrast between what is possible in labs and in demonstration projects in supportive environments, and what is possible in the real world.

McKinsey Global Institute (2018 page 8) notes that adoption of AI has some particular barriers at present. These are a lack of training data, difficulties in getting access to large enough datasets (especially in healthcare), difficulty in explaining results to human decision-makers, difficulties in generalising from one model to another which means each solution is only an answer to a relatively specific trouble, and the risk of bias in training data that leads to correspondingly biased outputs.

[7] https://ai.googleblog.com/2018/11/improved-grading-of-prostate-cancer.html

[8] https://jamanetwork.com/journals/jama/fullarticle/2588763

[9] https://blogs.microsoft.com/ai/machine-translation-news-test-set-human-parity/

There is also no shortage of popular disgruntlement with cognitive tasks that are trivial for humans being done slightly imperfectly by computers (like scheduling meetings),[10] and evidence that the hype around new labour-saving or job-killing innovations has got out of sync with what has actually been achieved (as in the case of self-driving cars).[11]

Lee (2018) adds a set of barriers specific to robots when operating outside of controlled factory or warehouse environments. He says the current state of AI technology is sufficient to programme adult cognitive functions, but creating a robot that can perform physical tasks even as well as a toddler is beyond human capability just now. AlphaGo, an AI programme, beat Lee Sedol the world's best Go player, but it relied on a human to put its pieces on the board.[12] Flippy the burger flipping robot struggled to integrate with its human co-workers and now has no allocated shifts at the fast food joint where it briefly worked,[13] and Laundroid, a $16,000 laundry-folding machine takes five or ten minutes to fold a t-shirt and struggled in a public test to do anything with a garment that was not one of the demonstration models it was trained on.[14]

Of course, the set of tasks that is technically possible for machines continues to grow over time. Lee (2018) says that at present they can do work that requires minimal social interaction and either limited dexterity in a structured environment, or limited creativity. The types of tasks where humans are at present expected to have enduring advantages over computers include those involving creativity, initiative, leadership or helping other people, or physical work in unstructured environments.

Even if the capabilities of AI never improved, existing capabilities are far beyond what is presently deployed. So we need some explanation of why organisations are not already taking advantage of the new possibilities presented by AI. What is technically possible represents the limit of adoption for AI, but there are three significant constraints.

What is permissible

Adoption of technology is limited by its "social licence", the ongoing approval of the operating practices of organisations that is necessary to operate in a society. A social licence is a general concept that has legal, ethical and cultural elements. It is broader than mere legal compliance because use of particular technologies typically will not require an actual licence. And it is broader than business competition because societal permission does not rest on customer takeup.

OECD (2017, page 278) notes the importance of a social licence to the development and deployment of technology:

> *"Strong public concerns can shape the direction, pace and diffusion of innovation, and even block its progress. This is the case even where technical and economic feasibility have been demonstrated, the rationale for adoption appears sound, and large investments have been undertaken."*

It gives the example of GMOs in Europe where negative public sentiment has resulted in "lower funding levels, high regulatory rejection rates, and lower levels of innovation than in other jurisdictions" (page 278).

The social licence provides a permissive space within which organisations can operate. The available social licence depends on the characteristics and potential impacts of technology. "Technologies that are perceived to be irreversible, out of human control, and/or capable of catastrophic failure" as well as technologies are novel and less well-known, outside of human perception and delayed in their manifestation of harm are likely to prove more problematic (page 283).

It is not just a question of education about technology but also of values. OECD (2017, page 283) notes that "where deeply held values and personal identities are at stake, science-based accounts [of technology impacts] are dismissed even by the most [science] literate". And that "value conflicts, distributional concerns and failures of trust" in institutions that manage risk are more important than a lack of knowledge or understanding.

---

[10] https://www.wired.com/story/xai-meeting-ai-chatbot/

[11] https://arstechnica.com/cars/2018/12/uber-tesla-and-waymo-all-struggled-with-self-driving-in-2018/

[12] https://www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611/

[13] https://www.theverge.com/2018/3/8/17095730/robot-burger-flipping-fast-food-caliburger-miso-robotics-flippy

[14] https://www.theverge.com/2018/1/10/16865506/laundroid-laundry-folding-machine-foldimate-ces-2018

Differences in technology characteristics and presentation might see particular innovations welcomed (as for example RocketLab in New Zealand),[15] taxed (as for AirBnB in Auckland),[16] delayed or outlawed (as for Uber in numerous places),[17] or otherwise controlled (as with the use of equipment made by Chinese firm Huawei in 5G telecommunications networks in New Zealand, Australia and the United States).[18] It might take time for an official response to emerge (see e-scooters in Auckland, Christchurch, and Dunedin at present),[19] and there is no reason to think that reactions will be the same for AI in general, rather than for particular services that use AI technologies. For example, preferences for human interaction may preclude the automated of some occupations, like care for the elderly and children (Borland and Coelli, 2017).

The swirl of narrative about big impacts on jobs and incomes makes us think that there will be continued high public interest in the adoption and use of AI technologies, especially those that might substitute for work that is presently done by humans.

General sentiment about developments in science technology in New Zealand is very positive, with interest in how science and technology can improve society and the state of the environment amongst the major motivations (Nielsen 2017). Shoham (2018, page 43) cites evidence that 30 per cent of popular media articles that use the term "artificial intelligence" are now positive in tone, up from about 12 per cent in 2016.

On the other hand, OECD (2017) cites a 2016 British poll that one third of respondents believe that the rise of AI is a "threat to humanity" and cautions that technology needs acceptance not just in general, but also from consumers and investors and from the populations that are affected by its deployment.

The OECD recommends that those seeking a social licence should take into account social goals and concerns in their work as far as possible and as early as possible. For AI in particular, the OECD suggests early funding of research into the "human and social dimensions of AI technology" integrated alongside technical research. One helpful approach could be to steer the technology in directions that respond to social concerns and serve wider goals rather than just economic ends, although the large-scale use of AI for social good also introduces new risks (McKinsey Global Institute, 2018b).

Some have expressed concern that seeking social licence for AI is impossible ex ante because of the wide span of AI impacts (OECD 2017, page 283). It is also not clear who should do the work of understanding stakeholders' views and establishing societal assent. The general purpose nature of AI technology means that there are few institutions with spans wide enough to encompass communication and learning and public engagement on all the issues.

The OECD is making some efforts in this direction at a high level. Its Committee on Digital Economy Policy has established an Expert Group on AI in Society to devise some principles to go to the OECD Council for endorsement in 2019 that would foster trust in and adoption of AI. The intention is to create some guidelines for countries on how to create national AI strategies that ensure that AI technologies benefit and empower as many people as possible.[20]

There is also a lot of attention being paid to the ethical issues around the development and use of AI. As just three examples among many, the Partnership on AI is a global multi-stakeholder body set up in 2016 to build understanding and formulate best practices on AI,[21] the IEEE has a standards body working on AI ethics,[22] and Google has published a set of internal AI principles that it uses to control its work.[23]

15 https://www.nbr.co.nz/article/looming-rocket-lab-launch-spurs-govt-set-space-regulatory-regime-b-190324

16 https://www.nzherald.co.nz/business/news/article.cfm?c_id=3&objectid=12178787

17 https://www.independent.co.uk/travel/news-and-advice/uber-ban-countries-where-world-taxi-app-europe-taxi-us-states-china-asia-legal-a7707436.html

18 https://www.stuff.co.nz/business/industries/108950894/nzs-big-call-on-huawei--politics-best-explains-5g-ban

19 https://www.stuff.co.nz/national/107117537/end-of-the-lime-auckland-councils-scooter-trial-set-to-end-on-january-14

20 http://www.oecd.org/going-digital/ai/oecd-initiatives-on-ai.htm

21 https://www.partnershiponai.org/

22 https://standards.ieee.org/industry-connections/ec/autonomous-systems.html

23 https://www.blog.google/technology/ai/ai-principles/

One specific issue that has received a lot of attention is the transparency of algorithms used in the processing of data. The UK House of Commons began an inquiry in 2017 into the use of algorithms in decision-making, including the extent to which they exacerbate or reduce bias and how they might be controlled. It generated a set of recommendations for government agencies, including the government's newly established Centre for Data Ethics and Innovation (House of Commons 2018).

The New Zealand government undertook a similar exercise within government in 2018, reviewing in particular the use of operational algorithms that impact significantly on individuals or groups based on a survey of 14 government agencies. It concluded that humans review and decide on almost all significant decisions made by the agencies, and made a set of recommendations to improve transparency, increase collaboration, and prepare for a future with greater use of AI in agency processes (Internal Affairs and Stats NZ 2018).

There was public commentary to the effect that agencies should make more use of algorithmic approaches to decision support (Kenny 2018) and also that there were more ethical, legal and technical challenges even at the current level of use that remained to be resolved (University of Otago, 2018). On the other hand, Zerilli et al (2018) point out the risks of requiring a standard of transparency from computerised algorithms that is higher than that expected of human decision-makers. And Ito (2018) is a cautionary tale about the perils of moving too quickly to damn algorithmic approaches.

There are examples from the wider data ecosystem of direct efforts to define and secure a social licence. The Data Futures Partnership, an independent group appointed and funded by the government to improve the data sharing ecosystem in New Zealand, undertook a guided public discussion process on the use of integrated data. It developed eight questions for organisations that can be used to isolate how generally acceptable their data use is likely to be to those who have provided the data (DFP 2017). The Social Investment Agency conducted a set of public engagements to find out what people and organisations thought about its "social investment" approach and the protection and use of personal information in the social sector (Internal Affairs and Stats NZ 2018, page 28).

Apart from social licences, the deployment of more sophisticated AI technologies could also mean more direct challenges to existing regulatory frameworks. One specific question that has been raised is the impact of New Zealand's copyright regime on the ease of deployment of machine learning solutions that rely on ready access to data (Deloitte Access Economics 2018).

What is profitable or effective

Even if AI technologies are technically possible and socially permissible, they will have limited adoption unless they lower firms' production costs, and provide more efficient ways for government and other not-for-profit organisations to progress towards their goals.

There are certainly an increasing number of useful practical applications of AI. AI Forum (2018, page 35) reports New Zealand evidence from a survey of early adopting organisations. The five most attractive features of AI solutions were in making sense of large volumes of data, automating work, supporting decision-making, reducing costs with automation, and improving business processes.

There is also some evidence saying AI solutions will boost firm productivity. McKinsey Global Institute research (2018 page 4) says that firms that fully absorb AI tools into their enterprises in the next five to seven years could double their cash flow relative to a scenario where they do not invest in AI. On the other hand, firms that do not adopt AI technologies at all or that have not fully absorbed them by 2030 could suffer a 20 percent decline in their cash flows. Competition will shift market share from laggards to front-runners.

McKinsey Global Institute (2018, page 39) estimates that two or three percent of firms will be front-runners, another 20 to 30 per cent will be followers, and about 60 to 70 per cent will be laggards. Front-runners will be firms that have a higher propensity to invest in AI, positive views of the business case for AI, and existing strong use of digital tools. Extending on this last point, companies that are more digitally mature will have annual AI adoption and absorption up to 12 percentage points higher than firms that are less digitally mature.

That said, there are also specific costs associated with using AI technologies. There is a set of technical precursors to make AI readily useable and useful. These include accessible big data (especially for machine learning), capabilities in data science, and some AI-specific methods, platforms and applications. The maturity of data management and governance within organisations is super important. McKinsey Global Institute (2018 page 25) says that even if the basic technological systems are in place, companies cannot generate value from AI without the skilled workers and experience necessary to take advantage of the technology and mobilise change within organisations. This makes AI technologies more likely to be adopted by larger organisations.

There has been some investigation of the readiness of government agencies to make use of AI solutions. A 2017 analysis by Oxford Insights ranked New Zealand in 9th position among the 35 OECD countries for AI Readiness in government. Its performance could improve with a stronger score in innovation (linked to lifting low R&D spending), a larger number of AI startups, and by improving what is a relatively limited availability of data.[24]

Timing

Like all technologies, the impacts of AI depend on the extent to which it is built in to organisations' production processes (Varian et al 2002, page 16). Even if AI technologies are technically possible, socially permissible and profitable, it will take organisations time to fully integrate them. This is a long process.

To illustrate how this works in more detail, we can use an example from a previous technology revolution, the adoption of Internet services by businesses. A stylised adoption cycle for bringing AI into an existing business process has five steps:

FIGURE 2

A stylised adoption cycle

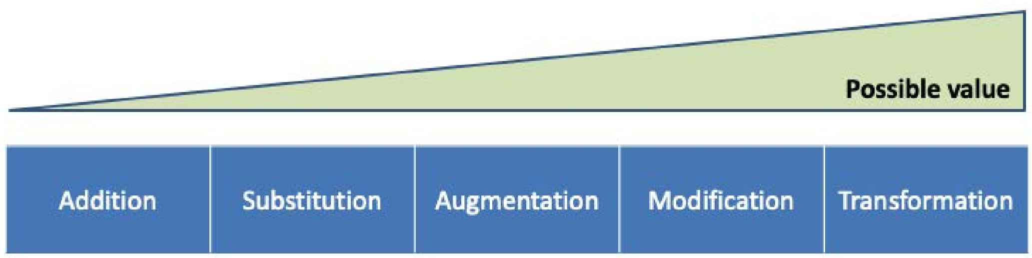

- Addition – An organisation introduces a new technology for some business purpose. It might continue to rely on its previous method in this early stage and run both systems in parallel. For example, we spoke to a medium-sized tourism operator in 2014 who had a web-accessible booking platform but still kept a paper-based system as the official record (Sapere 2014, page 15)
- Substitution – The organisation replaces some part of its process with the new capability. In our study, email had (mostly) replaced fax in tourism.
- Augmentation – The new technology enables new capabilities that the firm integrates into its process. An online booking system makes it easier for customers to find out about accommodation availability and to book themselves, even when the office is closed.
- Modification – The new technology changes the previous process. An online booking system enables a tourism operator to find customers and take bookings through third party engines. Those third parties can become important purchase pathways for customers and increase competition for tourism operators that used to rely entirely on a direct sales channel.
- Transformation – The new technology replaces the previous process. The tourism operator can get out of the business of running a booking system entirely, and just focus on raising awareness and on delivery of its tourism service.

How this cycle will look with AI is unclear at this point. But it does seem likely to take a long time. In previous work, we estimated that it could take 10 to 15 years for mass business adoption of AI, and then perhaps as long as another 25 years to get to transformation.

[24] https://www.oxfordinsights.com/government-ai-readiness-index/

McKinsey says that by 2030 up to 70 percent of firms may have adopted at least one type of AI technology, but less than half will have fully absorbed all five categories of AI that were included in the analysis.

This would be relatively rapid ("at the high end of what has been observed with other technologies") because of the number of ways in which AI can be used, and because returns from early adoption seem so large (McKinsey Global Institute 2018, page 25). Lee (2018) agrees that AI will be faster to be adopted than previous technologies because diffusion of digital algorithms is easier (needing just a computer and Internet connectivity), because funding for innovative enterprises is more readily available, and because the innovation effort is widely distributed across countries.

What can be said about adoption of AI now and in the future

Adoption could be measured as the proportion of all organisations that are using AI solutions in much the same way as one measures broadband adoption by businesses. Note that this means the proportion of organisations that are actively deploying AI technology. There will also be many users of services provided by others that are AI-powered without knowing that they are benefitting from AI.

As yet, it is too soon for there to be formal comparable international data on the use of AI.

What information there is at present on adoption relies on private survey evidence. AI Forum (2018, page 35) reports an adoption rate of 20 per cent in New Zealand based on a small survey of firms in technically advanced sectors, and described AI in New Zealand government agencies as being "disconnected and sparsely deployed".

McKinsey Global Institute (2018 pages 24, 51) reports that about a third of firms are using AI to some extent, based on various surveys including a survey of firms in ten large developed digitally advanced economies. It also analyses indicators of AI investment, research and the potential productivity boost and enablers like digital absorption, labour market structure and national connectedness across 41 countries and divides countries into four groups. The United States and China form the first group of active global leaders in the supply of AI with strengths that set them apart from all other countries. New Zealand is in the second group of countries with other small globally connected economies and some large countries with the ability to generate significant benefits from AI. The third and fourth groups are those with weaker economic and digital starting positions.

McKinsey's modelling says that economies with higher readiness to benefit from AI could reach absorption levels 23 percentage points higher than those of slow adopters by 2030 (McKinsey Global Institute 2018, page 34)

In McKinsey's assessment, New Zealand has an above average productivity boost available from automation because of relatively high cost labour. New Zealand has a well-educated population, and a flexible labour market structure. It rates as mediocre on the level of AI research, digital absorption, on innovation foundations like R&D investments, and on digital connectedness.

New Zealand does have some pre-cursors for AI adoption well-established. Investment in ICT by firms is the highest in the OECD, and broadband and computer takeup are universal at least amongst businesses with six or more employees where this is measured. New Zealanders are skilled at making use of technology in problem solving on OECD data. This suggests that New Zealand workers might find it easier to take advantage of new technological tools than workers in other countries.

On the other hand, the AI Forum survey reports low awareness of AI technology at present, even among technologically advanced firms. While many New Zealand decision makers were aware that AI will have transformative effects in general, 78 per cent of organisations reported that a major barrier was that parts of their business do not understand AI or its potential.

This might be part of a more general feature of the economy, and reflected in the fact that New Zealand is typically top 20 but not top 10 in global indices of innovation and competitiveness.[25] The New Zealand Productivity Commission says that it is rational for firms to be slower to adopt ICT in New Zealand, given that the domestic market is small, competition is limited by our distance from other countries, and ICT investments involve fixed costs (NZPC 2014, page 107). Paul Conway, the Director of Economics and Research at the

[25] For examples, New Zealand is ranked 13th in the World Economic Forum Global Competitiveness Index 2017-18, 22nd on the INSEAD/WIPO Global Innovation Index 2018, and 17th on the World Economic Forum Network Readiness Index 2016,

Commission has also fingered the predominance of cooperative business models and partly-privatised government agencies amongst New Zealand's largest entities as being a handbrake on innovation (Conway 2018, page 50).

McKinsey says that competition is the most important driver for technology adoption from their survey evidence ("the presence of rivals investing in AI plays a significant part in any decision by a company to invest" 2018, page 26). It reckons that competitive pressure would increase the absorption level of AI in a country by about 13 percentage points in 2030, ie, in a more competitive country full deployment of AI might reach 48 per cent of firms from only 35 per cent in less competitive economies. We would expect New Zealand to be at the bottom end of this range therefore.

IMPACTS

In this section, we outline possible economic and social impacts from the adoption of AI, largely drawn from our previous work on these issues for the AI Forum (AI Forum 2018) and from the most recent international studies.

As ever when discussing the impacts of technology, the main effects are not driven by the tech sector itself, but from organisations outside the tech sector taking advantage of technology in what they do (Oulton 2010). And also note that it takes a long time for the effects of technology to show up in the official statistics. Economists are still debating the productivity impacts of computerisation (eg Brynjolfsson and McAfee 2014, Gordon 2016).

Aggregate economic measures

AI is expected to be a big deal for productivity and economic growth. There are several recent studies of the impacts. These have different assumption about adoption built in:

- Accenture (2016, 2017) says that AI could double growth rates in 12 developed countries by 2035, and boost labour productivity by up to 40 per cent. especially in the tech sector, financial services and manufacturing.
- McKinsey Global Institute (2018) reckons AI could deliver additional economic output of around $13 trillion by 2030, boosting global GDP by about 1.2 percent per year. It predicts a 26 per cent gross increase in GDP less ten per cent for the costs of the change and the investment required. It notes that the large initial investments and ongoing refinement and transition costs might limit adoption by smaller firms.
- In our previous work with Schiff Consulting (AI Forum 2018) we calculated that use of AI technologies could boost New Zealand GDP by up to $54 billion by 2035 across 18 industries that account for 96 per cent of output.

McKinsey Global Institute (2018, page 6) observes that at these levels the impacts of AI would compare well with other general-purpose technologies throughout history, with impacts about four times as big as the introduction of steam engines during the 1800s (labour productivity boost estimated at 0.3 per cent per year), three times as big as industrial robots during the 1990s (0.4 per cent), and twice the impacts of the spread of IT during the 2000s (0.6 per cent).

McKinsey Global Institute (2018, page 38) says the digital maturity of sectors correlates with their level of AI adoption and absorption. It therefore reckons that the impact of AI in the telecommunications and high-tech sector could be twice as big as in healthcare.

Total number of jobs

These aggregate economic impacts generally come about through the impacts of AI on the demand for labour, and from the ability to generate new products and services. Broadly the labour market impacts are expected at this point to be similar to previous technology waves:

- Total employment is expected to be unaffected (though opinions vary),
- Some occupations will shrink while others will grow,
- A small number of new jobs will be created, and
- There will be very substantial change in existing work.

Sober predictions are that overall job growth will be in line with historic trends (Borland and Coelli 2017, Autor 2015, EOP 2016b, OECD 2017, Meyers and Besanko 2017, PWC 2018). For example, Borland and Coelli (2017) conclude that "the aggregate amount of work available to the Australian population on a per capita basis has remained stable" since the mid-1960s, measured as annual hours worked per capita. They say that there is no evidence that computers have increased the pace of destruction of routine jobs through the 2000s. In our work with Schiff Consulting, we concluded AI-driven job displacement would account for only 10 per cent of normal job creation and destruction over the next 40 years.

McKinsey Global Institute (2018, page 45) is a little more concerned, saying that total pressure on full-time employment will be "more limited than many fear", leading to an average global scenario where total full-time-equivalent employment will remain flat compared with today. Automation could, on average, substitute for around 15 percent of existing time worked globally by 2030 (page 14).

And there are some commentators who are more alarmed about the aggregate employment impacts. For example, technology investor Lee (2018) thinks that AI will reduce employment in the United States by nearly half: 38 per cent following estimates from PWC (2018) and an additional 10 per cent as a result of entire industries being replaced by new approaches to production enabled by AI. He says that, in comparison with other general purpose technologies like steam power and electricity, AI can be expected to boost productivity but not employment. Taking into account AI-related job growth plus inertia in the takeup of AI, he reckons that unemployment in the United States will "only" be 10 to 20 per cent, comparable with the Great Depression of the 1930s but on a long-term or permanent basis.

We are on more settled ground predicting the decline of some occupations and the rise of others.

- As mentioned above, Frey and Osborne (2013) and Artnz et al (2016) are the most popular references, predicting the demise of about half or nine per cent of existing jobs in the United States to AI respectively, depending on whether one makes one's predictions based on entire jobs or on particular tasks.
- Elliot (2017) says that there will be strong economic pressure to use computers to deliver tasks that involve basic literacy, numeracy and reasoning skills across the economy, and that this will reduce the number of workers using low and mid-level literary skills in particular.
- Bain and Company (2018) think that automation could eliminate 20 to 25 per cent of current middle to low-income jobs by the end of the 2020s.
- In line with the technical capabilities of AI, routine jobs with limited social interaction are expected to be most at risk. Lee (2018) reckons that service jobs are more protected than routine computational ones because robots are less capable physically than AI algorithms are cognitively. He therefore expects AI impacts to hit white collar workers first.

Direct job growth as a result of AI is expected to be relatively modest. McKinsey Global Institute (2018, page 14) says that between 1980 and 2000 in the United States, about 4 to 9 percent of the workforce were employed in job categories that did not exist 10 to 15 years earlier. Making some adjustments for the characteristics of AI, they conclude that new jobs driven by investment in AI could contribute about five percent to employment by 2030 (page 46).

The indirect job-creating impacts of AI are expected to be much bigger. For instance, McKinsey Global Institute (2018, page 46) estimates that the advent of the personal computer has enabled the creation of nearly 16 million net new jobs in the United States since 1980, and 90 per cent of these jobs are outside the technology sectors.

Change in jobs

The big impact of AI is change in the way that people do their jobs, similar to the use of ICT in the workplace since the 1980s.

One way this has been measured is as the proportion of tasks within a job that can be automated. For example, McKinsey Global Institute (2018, page 21) divided work activities into seven types and explored the technical potential for automation for each type. Three of the seven types have about 60 to 80 percent potential for automation: "performing physical activity and operating machinery in predictable environments", "processing data", and "collecting data". The other four are much less automatable, at around 10 to 25 percent: "performing physical activities and operating machinery in unpredictable environments", "interfacing with stakeholders", "applying expertise to decision making, planning, and creative tasks" and "managing and developing people".

AI is expected to shift the jobs mix toward tasks requiring high digital skills and those involving non-repetitive tasks. More of us will be working with machines and directly with other humans. McKinsey Global Institute (2018, page 42) says that the share of employment made up by jobs with repetitive activities and requiring low digital skills could be reduced by 25 per cent, while jobs involving non-repetitive activities and those that require high digital skills could rise from 40 per cent to more than 50 per cent.

Education is protective. OECD (2017) says that fewer than five per cent of workers with a tertiary degree are at a high risk of losing their job due to automation, compared to 40 per cent of workers with a lower secondary degree. This is good news for women, since more women than men are now tertiary graduates across the OECD.

McKinsey Global Institute research says that up to 14 percent of workers might need to change occupations. While some may change roles within the same company, others may need to move to new sectors or cities. OECD (2018) says that the workers most at risk are those that are least connected to the education system and that for this reason existing training systems will not reduce the risks for affected workers.

There would be a set of flow-on effects economically and socially from big changes in employment. McKinsey Global Institute (2018, page 4) sees greater competition for workers that are able to develop and use AI tools, and "structural excess supply" for the high portion of people lacking the digital and cognitive skills needed to work with machines.

Social impacts

There is a wide range of potential applications for AI technologies outside of the commercial sector, including in social services and delivery by government agencies.

As just a handful of examples from recent reports, AI is being explored to more readily study the health of native bird populations,[26] more quickly diagnose sick cows on dairy farms,[27] improve the quality of bail decisions in the criminal justice system (Kleinberg et al 2017), or to detect more counterfeit goods crossing borders.[28]

McKinsey Global Institute (2018b page 3) sees possibilities for large impacts in health, education, security and justice, and social equality and inclusion. It notes for example in the health sector:

> *"AI-enabled wearable devices, which can already detect potential early signs of diabetes through heart rate sensor data with 85 percent accuracy, could potentially contribute to helping more than 400 million people afflicted by the disease worldwide if made sufficiently affordable."*

but also sees a range of barriers to be overcome, especially in access to data and people with the skills to make use of it. Useful data might be privately collected and only available commercially or collected by government agencies but still difficult to access by agencies outside of government.

RESPONSES

We now turn to look at what potential responses to the opportunities and challenges presented by AI might be.

As ever with technology adoption, the main action is on businesses, government agencies and other organisations to take advantage of AI solutions where that makes sense for them. Some governments around the world have also taken specific policy actions under a documented national AI strategy.

26 https://www.nzherald.co.nz/nz/news/article.cfm?c_id=1&objectid=11923979

27 https://www.computerworld.co.nz/article/620747/ai-diagnose-sick-cows/

28 https://www.scmp.com/news/hong-kong/hong-kong-law-and-crime/article/2158400/new-supercomputer-helps-hong-kong-customs

In this section, we present a taxonomy of policy options and examples from other countries or other technologies. We start by looking at why organisations or governments might choose to do anything at all at this early stage.

Why do anything

As ever, the private incentives for firms, government agencies or not-for-profits to act depend on their view of the benefits available from AI technologies as compared with the effort involved in deploying them. The numbers from McKinsey Global Institute (2018) say there are big competitive advantages from being an early mover, because coming later to the party means a disproportionately smaller slice of cake. But the pre-cursors for making effective organisational use of AI are costly to acquire, which will mean that larger and more technically advanced firms will be the front-runners.

There are specific roles for government in setting broad technology directions for its agencies, allocating research funds, helping private firms to coordinate their efforts, and lifting the profile of firms that are making smart use of technology, including AI. It is hard to see a role for government in pushing specific technology investments by businesses, but it does have a range of existing interventions to encourage firms in general to make use of technology that could readily by used to lift awareness and boost adoption of AI as well.

The OECD reckons that policy matters for adoption of AI and for its impacts (2017, page 294):

> *"national investments and strategies can and will exert a profound influence on the direction of technological change"*

One challenge McKinsey sees is that adoption of AI could widen wealth gaps between countries, with those countries that establish themselves as AI leaders capturing twice as much economic upside as countries that do not (McKinsey Global Institute 2018, page 3). This would be an argument to move early. Lee (2018) says that there is little point. Countries that are global leaders will need vibrant entrepreneurial ecosystems, access to talent, research funding, a large base of users and effective venture capital networks. His view is that only the United States and China will have all these factors.

There are also arguments about so-called "superstar" effects, where there impact of having more customers and more data to hand gives an increasing advantage to larger players in the development and deployment of AI. Depending on one's view of the future, this could lead to widespread monopolisation of markets, and big inequalities in performance between large countries that effectively deploy AI and everybody else. Lee (2018), for example, foresees a global economy where low cost labour gives no economic advantage.

Response options

By November 2018, eighteen countries had released some form of national AI strategy (Dutton 2018). The first was Canada, in March 2017. Another ten governments indicated the intention to present something in the coming year (including New Zealand). Half of the 18 strategies outline specific policies and include funding, and the other nine present objectives to guide future decision-making.

Dutton (2018) says that the strategies vary substantially in their focus from country to country, but are made up of efforts in eight areas of public policy. He isolates four types of strategies: those that focus on AI research and building talent, those that aim to encourage private sector AI adoption, comprehensive strategies that cover all or most of the eight areas, and documents that provide guiding principles for future policy only.

On his analysis:

- Encouraging private sector AI adoption is a priority for eight of the 18 of the national strategies
- Scientific research and building AI talent are priorities for seven countries.
- Work on ethical AI standards, broader data and digital infrastructure developments or to expand the use of AI in government are priorities for a few countries.
- Making sure AI is inclusive, and skills development for the future of work are the least and second-least prioritised policy areas, respectively.

- Only five strategies touch on all eight policy areas, while seven are quite focused and touch on four or fewer areas.

Some specific examples from particular countries:

- Canada and South Korea both have research-oriented strategies involving establishing new AI research institutes and providing funding to attract, train and retain AI talent in the country.
- The UK strategy touches on all eight policy areas, lifting funding for AI research and education, encouraging private sector investment, creating a centre to explore issues in data ethics and a new AI Council to advise the government, and establishing new arrangements to make trusted data sharing easier.
- The United States has no Federal AI strategy as yet although there were two reports in 2016 from the previous administration (EOP 2016a, 2016b). AI efforts are largely privately led, with the Federal government investing in AI research and removing regulatory barriers. NSTC (2016) is a plan for research and development of AI that identifies seven priorities for Federally-funded AI research. The US House of Representatives has explored these issues and called on the Federal government to increase its engagement to maintain United States global leadership in AI and address its challenges (USHoR 2018). Carter (2018b) is a private attempt to define a strategy based on goals of promoting safe and responsible development, and maintaining United States technical leadership in what it calls "Machine Intelligence".
- China's Next Generation Artificial Intelligence Development Plan is a three-phase strategy to make China the world leader in AI theories, technologies, and applications by 2030. It includes efforts across all eight policy domains, including making breakthroughs in fundamental AI research, making China the world's primary AI innovation centre, and leading the world in standards setting and AI ethics (Dutton 2018).

The table below shows the amounts of funding announced for AI for the nine national strategies that include funding from Dutton (2018), with the rows sorted by country population to make the spending figures more comparable.

FIGURE 3

Announced investments in nine national AI strategies

Country	Investment ($USD)	Population	Comment
Denmark	20m	6m	
Singapore	92m	6m	Over five years
Taiwan	1.2b	24m	Over four years
Australia	22m	25m	
Canada	95m	37m	
South Korea	2b	51m	
United Kingdom	1.2b	66m	
France	1.75b	67m	Over five years
EU	1.75b	513m	

Based on the Dutton (2018) analysis and the work of the OECD,[29] a non-exhaustive menu of options for governments to respond to the opportunities and threats created by AI includes the following two lists.

The first set is for actions that we can see some countries are already embarked upon:

29 http://www.oecd.org/going-digital/ai/

- Talk widely about the opportunities and challenges of AI to encourage action and build understanding.
- Encourage adoption of AI technologies by businesses, government agencies and others to speed the achievement of benefits.
- Fund technology research to build national AI capability or to develop new AI-powered products and services for commercial advantage – Carter (2018) says that private sector research will be focused on near-term and commercially useful applications (like machine vision, natural language processing, and autonomous vehicles). His view is that governments should solve instead pressing long-term issues, like how to move beyond machine learning, and how to make AI decision support more transparent, more resilient and less data-intensive.
- Explore wider policy, regulatory or ethics issues that will help enable actions by others or that will be important in the longer-term – There will be opportunities to regulate AI under existing authorities. Self-regulation of some aspects of these new technologies seems unlikely to be sustainable. As Carter (2018) notes, "the tech sector's credibility on managing issues of privacy, security, and broader social impacts of technology is at an all-time low."
- Make more data that is collected by government agencies available for use in developing AI-powered products and services.
- Develop specific programmes to use AI in the provision of government services or policy advice.
- Make investments in training or greater resilience for workers expected to be affected in the move to the future of work – There has been a lot of analysis of the characteristics of who might be affected. Elliot (2017) says that the standard policy response of education might not be as effective this time because there is no education system in the world that prepares the vast majority of adults to perform better at literacy, numeracy and problem-solving skills than the level computers are close to operating at. He recommends instead helping workers to develop skills that the computers cannot replicate.
- Look to attract, retain or train people with relevant skills, through coordination with education providers and immigration settings.
- Make policy efforts to ensure that the AI community and the goals to which AI is devoted are inclusive and diverse.

The second set is for actions that have been discussed in relation to AI, or used in the case of other technologies but do not feature so far in national strategies:

- Tax automation or technology deployment to slow adoption and fund the social costs of change – A body in the EU regulatory apparatus talked about this option but did not adopt it in their final rule-making.[30] Generally speaking it would be odd to tax R&D (it is more normal to support it through the tax system) but more possible to tax specific applications of AI that are thought to have negative impacts for some workers or industries.
- Use funding to shape technology development to encourage solutions to policy issues or to minimise negative impacts – Prizes, national challenges, or defined focus areas for research could give some teeth to the idea of influencing technology development or deployment in directions that are less likely to make human labour redundant, for example. They could also be used to encourage the application of technology to big societal challenges that do not have an obvious economic payoff. Examples might include reducing road deaths, the prison population, or the spread of invasive species.
- Directly limit technology choices to control the path of development or favour particular firms or technologies – The limits on the use of Huawei communications equipment mentioned above is a recent example of this approach. A set of geostrategic issues come into play in commercial rivalry where countries see national security issues or fundamental questions of comparative advantage. There are also broader motivations for this type of policy approach, for example the campaign to ban the development, production and use of fully-autonomous weapons.[31]

No matter what path is chosen, it is always helpful to keep an eye on what other countries are doing. The OECD now usefully provides a clearing house for AI-related developments around the world. It also plans to launch in

30 https://www.reuters.com/article/us-europe-robots-lawmaking/european-parliament-calls-for-robot-law-rejects-robot-tax-idUSKBN15V2KM

31 https://www.theguardian.com/technology/2018/apr/09/killer-robots-pressure-builds-for-ban-as-governments-meet

2019 an online AI Policy Observatory to bring together evidence, analysis and options on how to ensure the beneficial use of AI.[32]

As with other general purpose technologies, the widespread use of AI will raise broader policy questions, for instance in education systems, labour market matching, the regulation of data use, income support, competition policy, and industry structure, amongst others.

Lee (2018) says that the sensible policy responses are focused on the labour market, viz retraining workers, reducing work hours to share work more widely, or redistributing income to maintain standards of living. But he is sceptical that these can actually work in a systematic way or speedily enough. He calls instead for the fundamental revaluation of human-centred service jobs so that they pay more, and to ensure that humans work alongside AI to humanise the analytical perfection that AI might otherwise provide in too clinical a fashion.

First steps for policy

An early question for policy at this stage in the technology wave is to understand how many New Zealand organisations are using AI and how extensively. The best way to measure that is through surveys, the same way that the adoption of other technologies is measured. One can also look at the availability of relevant pre-cursors like appropriate data, people with the skills required, the number of jobs and incomes of those working in ICT occupations and in industries with higher uptake, and organisational awareness of the capabilities of AI. There is doubtless some existing public science funding relevant to AI. It would also be helpful to understand its extent and distribution.

It is likely too early in the cycle to see any significant economic or social impacts other than through case studies or examples of trials or deployments. But in due course these impacts will be able to be measured in the usual way, ie, economically through firm investments in technology and ultimately through growth in productivity, and in other fields through ongoing improvements in outcomes or efficiency.

References

Accenture (2016), "Why Artificial Intelligence is the future of growth"

Accenture (2017), "How AI boosts industry profits and innovation"

AI Forum (2018), “Artificial Intelligence: Shaping a Future New Zealand”

Andriole, Steve (2018), “Artificial Intelligence, China And The U.S. - How The U.S. Is Losing The Technology War”, Forbes, November 9, 2018

Arntz, Melanie, Terry Gregory and Ulrich Zierahn (2016), "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis", OECD Social, Employment and Migration Working Papers, No. 189

Autor, David (2015), “Why are there still so many jobs? The history and future of workplace automation”, Journal of Economic Perspectives, 29(3), 3-30

Bain and Company (2018), “Labor 2030: The Collision of Demographics, Automation and Inequality”, February 2018

Borland, Jeff and Michael Coelli (2017), “Will robots take our jobs?”, July 2017

Brynjolfsson, Erik, Andrew McAfee (2014), “The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies”

Carter, William A (2018), “Promising Start, but Few Details in House AI Report”, CSIS, September 26, 2018

Carter William A (2018b), “A National Machine Intelligence Strategy for the United States”, CSIS, March 2018

Conway, Paul (2018), “Can the Kiwi Fly? Achieving productivity lift-off in New Zealand”

32 http://www.oecd.org/going-digital/ai/oecd-initiatives-on-ai.htm

Data Futures Partnership (DFP) (2017), “A Path to Social Licence: Guidelines for Trusted Data Use”

Davis, Fred D (1989), “Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology”

Deloitte Access Economics (2018), “Copyright in the digital age: An economic assessment of fair use in New Zealand”, Report for Google, March 2018

Dutton, Tim, Brent Barron, Gaga Boskovic (2018), “Building an AI World: Report on National and Regional AI Strategies”, CIFAR

Elliot, SW (2017), “Computers and the Future of Skill Demand”, OECD Publishing, Paris

Executive Office of the President (EOP) (2016), “Preparing for the Future of Artificial Intelligence”, National Science and Technology Council Committee on Technology, October 2016

Executive Office of the President (EOP) (2016b), “Artificial Intelligence, Automation, and the Economy”, Executive Office of the President, December 2016

Frey, Carl Benedikt and Michael A Osborne (2013), “The Future of Employment: How susceptible are jobs to computerisation?”

Gordon, Robert J (2016), “The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War”

Gregory, Terry, Anna Salomons and Ulrich Zierahn (2016), “Racing with or against the machine? Evidence from Europe”, ZEW, Discussion paper no.16-053

Hanson, Robin (2001), “Economic growth given machine intelligence”, Technical Report, University of California, Berkeley.

House of Commons (2018), “Algorithms in decision-making”, Science and Technology Committee, Fourth Report of Session 2017-2019, May 2018

Internal Affairs and Stats NZ (2018), “Algorithm Assessment Report”, October 2018

Kenny, Katie (2018), “Algorithms are everywhere but the public sector seems scared to use them”, Stuff, November 13 2018

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, Sendhil Mullainathan (2017), “Human Decisions and Machine Predictions”, NBER Working Paper No. 23180, February 2017

Ito, Joy (2018), “What the Boston School Bus Schedule Can Teach Us About AI”, Wired, 11 May 2018

Lee, Kai-Fu (2018), “AI Superpowers: China, Silicon Valley and the New World Order”

Lunceford, Brett (2009), “Reconsidering Technology Adoption and Resistance: Observations of a Semi-Luddite”

McKinsey Global Institute (2018), “Notes from the AI Frontier: Modeling the Impact of AI on the World Economy”, September 2018

McKinsey Global Institute (2018b), “Notes from the AI Frontier: Applying AI for Social Good”, December 2018

Meyers, Max and David Besanko (2017), “Adjusting to Automation: Public Policy and the ‘Future of Work’”

National Science and Technology Council (NSTC) (2016), “The National Artificial Intelligence Research and Development Strategic Plan” Networking and Information Technology Research and Development Subcommittee, October 2016

Nedelkoska, L and G Quintini (2018), “Automation, skills use and training”, OECD Social, Employment and Migration Working Papers, No. 202

Neilson (2017), “Report on Public Engagement with Science & Technology”, Prepared for MBIE, November 2017

New Zealand Productivity Commission (NZPC) (2014), “Boosting productivity in the services sector: 2nd Interim Report Competition and ICT topics”

OECD (2017), “The Next Production Revolution: Implications for Governments and Business”

OECD (2017b), “Going Digital: The Future of Work for Women”, Policy Brief on The Future of Work

OECD (2018), “Putting faces to the jobs at risk of automation”, Policy Brief on the Future of Work

Oulton (2010), “Long term implications of the ICT revolution: applying the lessons of growth theory and growth accounting”

Parker, Emily (2018), “How two AI superpowers — the U.S. and China — battle for supremacy in the field”, Washington Post, November 2, 2018

PWC (2018), “Will Robots Really Steal our Jobs: An international analysis of the potential long term impact of automation”

Sapere (2014), “The Value of Internet Services to New Zealand Businesses”, Report for Google and Internet NZ, March 2014

Shoham, Yoav, Raymond Perrault, Erik Brynjolfsson, Jack Clark, James Manyika, Juan Carlos Niebles, Terah Lyons, John Etchemendy, Barbara Grosz and Zoe Bauer, (2018) "The AI Index 2018 Annual Report”, AI Index Steering Committee, Human-Centered AI Initiative, Stanford University, Stanford, CA, December 2018

United States House of Representatives (USHoR) (2018), “Rise of the Machines

Artificial Intelligence and its Growing Impact on U.S. Policy”, Subcommittee on Information Technology, Committee on Oversight and Government Reform, September 2018

University of Otago (2018), “Stock-take of AI in government a great start”, 26 October 2018

Varian, Hal, Robert E. Litan, Andrew Elder, Jay Shutter (2002), “The Net Impact Study: The Projected Economic Benefits of the Internet In the United States, United Kingdom, France and Germany”, January 2002

Venkatesh, Viswanath and Fred D Davis (2000), “A Theoretical Extension of the Technology Acceptance Model: Four Longitudinal Field Studies”

Zerilli, John, Alistair Knott, James Maclaurin, Colin Gavaghan (2018), “Transparency in Algorithmic and Human Decision-Making: Is There a Double Standard?”

This literature review was built on research conducted by Hayden Glass of Sapere (an independent economic consultancy), on behalf of AIFNZ.
Contact: Hayden Glass
hglass@srgexpert.com
Ph +64 21 689 176

The Artificial Intelligence Forum of New Zealand is a non-government association with a mission to harness the potential of Artificial Intelligence (AI) to help bring about a prosperous and inclusive Future New Zealand.

The rapid development of AI technologies presents major opportunities and challenges for our country: from creating world leading AI businesses, nurturing a pool of talented AI engineers, applying AI technologies to our agriculture, government, manufacturing and service industries to holding a meaningful national debate on the broader implications for society, New Zealand needs to actively engage with AI now in order to secure our future prosperity.

The Forum brings together citizens, business, academia and the government connecting, promoting and advancing the AI ecosystem to help ensure a prosperous New Zealand.

AI Forum of New Zealand
L1 Building C, 14-22 Triton Drive, Auckland 0632, New Zealand
Ph +64 9 394 7693
www.aiforum.org.nz