Horizon Scanning Series

The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing

Employment and the Workforce

This input paper was prepared by Alexander Lynch on behalf of Google Australia

Suggested Citation

Lynch, A (on behalf of Google Australia) (2018). Employment and the Workforce. Input paper on behalf of Google Australia for the Horizon Scanning Project “The Effective and Ethical Development of Artificial Intelligence: An Opportunity to Improve Our Wellbeing” on behalf of the Australian Council of Learned Academies, www.acola.org.

Google Australia Pty Ltd
Level 5, 48 Pirrama Road
Pyrmont NSW 2009

Tel: 02 9374-4000
Fax: 02 9374-4001
www.google.com.au

Precis

Organisations around the world, including governments, companies, academic institutions, non-governmental and intergovernmental organisations, are investigating how to best make use of artificial intelligence (AI) for human good, and are aligning on ethical principles and governance standards regarding the use of this technology. Australia is no different. But when it comes to upskilling and investment in AI education and AI implementation Australia lags global peers, and if left unresolved this will have substantial economic and geostrategic implications for Australia in the years ahead.

Artificial Intelligence and Australia

Although Australians are recognised as leaders in technology adoption as individuals, our institutions lag peer nations in rates of productivity enhancing technology investment, business sophistication, and knowledge and technology output levels.

In Australia, the proportion of publicly listed companies engaging in sustained productivity enhancing investment is around 9 per cent. This is close to one third the rate of the world leader, Switzerland, where productivity enhancing investments helped boost the productivity of their workers by at least 5 per cent between 2010 and 2015, and less than half the rate of publicly listed companies in the United States (20.3 per cent)[1].

The 2018 Global Innovation Index highlights the comparatively low sophistication of Australia's business sector and low knowledge and technology output levels as weaknesses, however, Australia has noted strengths that provide a relative advantage for AI adoption compared to nations in our region. These include strong human capital and research capabilities, effective government service delivery, and a sophisticated local market supported by robust infrastructure[2].

In scenarios where nations take an active role in workforce development through reskilling, the Economist Intelligence Unit found that Australia was the developed nation with the most to gain from machine learning, with compound annual growth (CAGR) in productivity increasing to 2.25 per cent against their baseline forecast of just 0.19 per cent[3].

Most Australians experience AI enabled services every day, and although Google Australia has a local engineering presence of more than 650 people employed in computer science

[1] *The Automation Advantage*, Alphabeta, 2017
[2] *Global Innovation Index 2018*, Cornell University, INSEAD, and the World Intellectual Property Organization, 2018
[3] *Risks and Rewards: Scenarios around the economic impact of machine learning*, Economist Intelligence Unit, 2017

roles, working on Google's AI-integrated products, in the case of many AI-enabled technology products in everyday use, the technology Australians rely on is built by developers in markets like the United States, United Kingdom and China, where governments have strategically invested in and supported computer science industry development. Advanced digital capabilities are deployed into Australia from data centres in other nations where there exists sufficient compute to run AI implementations that are already complex and rapidly increasing in compute intensity.

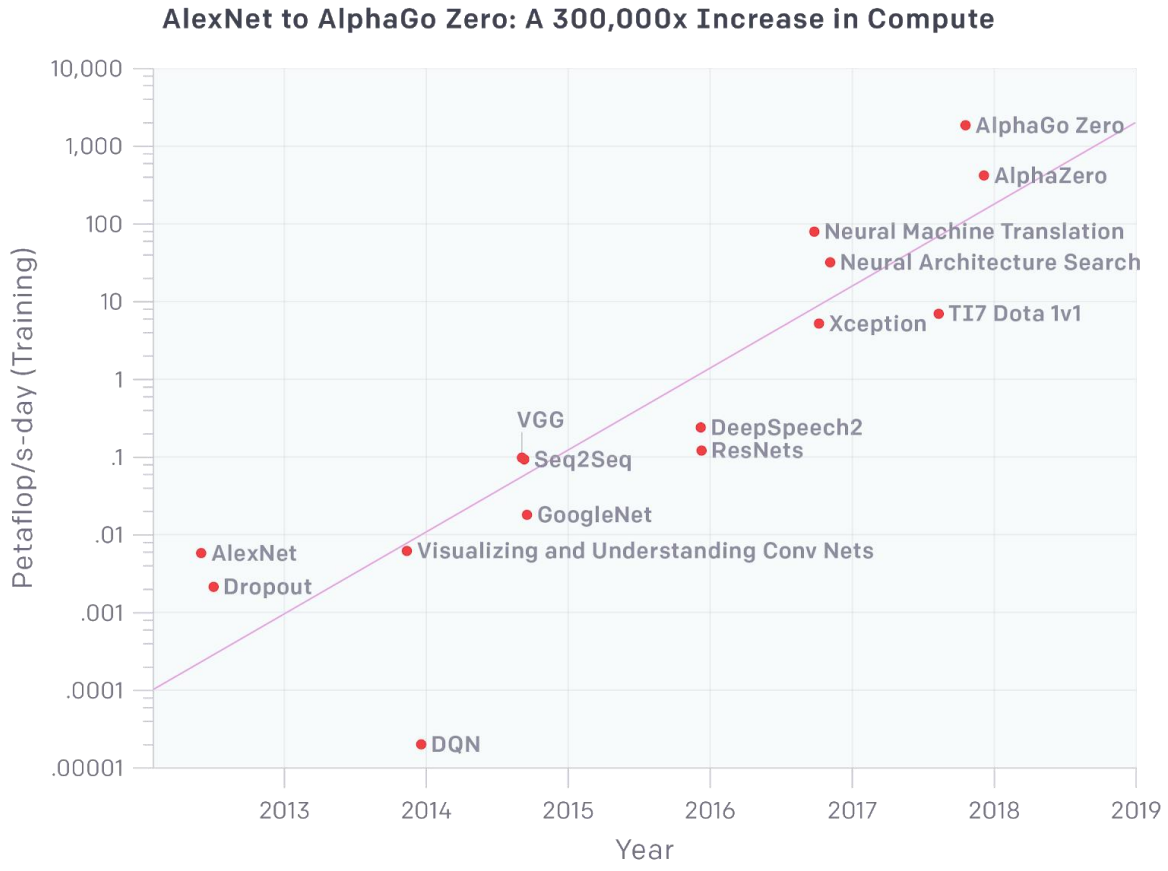

Source: Open AI

> *"Three factors drive the advance of AI: algorithmic innovation, data (which can be either supervised data or interactive environments), and the amount of compute available for training. Algorithmic innovation and data are difficult to track, but compute is unusually quantifiable, providing an opportunity to measure one input to AI progress. Of course, the use of massive compute sometimes just exposes the shortcomings of our current algorithms. But at least within many current domains, more compute seems to lead predictably to better performance, and is often complementary to algorithmic advances."*

Although we see strong demand from local organisations for guidance on potential use cases of AI, which as Australian organisations invest would over time would increase the rationale for investment in local compute infrastructure, there remain structural barriers to investment in Australia. For example, Australia lacks a Fair Use provision in copyright law to

facilitate the use of widely-distributed publicly available data for model training, presents an uncertain regulatory environment for digital technology, and outside specific and highly-capable organisations such as Data61, leaders across the economy exhibit low technical understanding of AI capability and the operation of digital systems more generally. This education gap can lead to failures of analysis including overgeneralisation and overestimation of AI use and capability in Australia.

More relevant for Australia's specific national challenge, nations worldwide are putting into place strategies to ensure they prioritise the development of AI capability, and leaders in the field are working to ensure they retain the long-term economic and geostrategic advantages that AI investment will provide. For example, in 2016 the United States National Science and Technology Council published *The National Artificial Intelligence Research and Development Strategic Plan*, which was followed in 2017 by the publication of *A Next Generation Artificial Intelligence Development Plan*[4] by China's State Council, and in 2018 the United Kingdom's *AI Sector Deal* from the Departments for Business, Energy & Industrial Strategy and Digital, Culture, Media & Sport, following their report *Growing the Artificial Intelligence Industry in the UK*[5].

Like any other tool, AI can both solve and create problems. Developing local capability and expertise will ensure Australia is capable of assessing and addressing the possibilities for the latter, and in the case of the former, is able to use AI to improve our institutions and society, and in doing so build the demand case for high-technology investment in Australia.

AI for Everyone

Given the potential AI has to help solve problems in a range of fields, Google views it as important that everyone has access to the opportunities this technology presents.

Google is the leading publisher of public machine learning research globally, and using the Category Normalized Citation Impact, a measure of how influential a research publication has been on academic discourse, the impact of research papers published publicly by Google is four to five times the world average.

Our commitment to open science has also led Google to develop and open-source the machine learning framework TensorFlow, which is now the largest machine learning community on GitHub. Researchers and developers around the world share deep learning models and datasets designed to make deep learning more accessible and accelerate machine learning research. We support the open source machine learning community with public, annotated machine learning datasets curated by Google engineers.

[4] An english language translation can be found here.

[5] A review of National AI Strategies published on Medium by the Founder and Editor-in-Chief of Politics + AI found that in the past fifteen months, Canada, China, Denmark, the EU Commission, Finland, France, India, Italy, Japan, Mexico, the Nordic-Baltic region, Singapore, South Korea, Sweden, Taiwan, the UAE, and the UK have all released strategies to promote the use and development of AI.

We have also partnered with Udacity to publish a free, public Intro to Deep Learning course, and have made a range of educational resources, including an online machine learning crash course, available for free.

AI supports human information acquisition through Google Search, navigation through Google Maps, and even allows us to provide these free, ad supported, public goods more energy efficiently by reducing our data centre energy consumption.

At a commercial level, machine learning systems created by Google can be accessed by third parties for their own use. For example, at the date of this submission these AI APIs include:

- Speech recognition to convert audio, either from a microphone or from a file, to text in more than 80 languages
- Machine translation between almost 100 different language pairs
- Natural language tools that allow analysis of the structure and meaning of text, and
- Image analysis through machine vision, including object detection and labelling, optical character recognition, and detection of explicit content.

Google has also made available AutoML, a service that allows third parties with limited machine learning expertise to create custom machine learning models.

Google has designed and deployed custom machine learning hardware, Tensor Processing Units (TPUs), specifically to support AI model training and implementation, and makes this technology available to machine learning researchers outside the company. A single networked pod of third generation TPUs sustains 100 petaflops of computing power. In comparison, Raijin, Australia's fastest supercomputer, located at Australia's National Computational Infrastructure National Facility in Canberra, delivers 3.8 petaflops at peak performance.

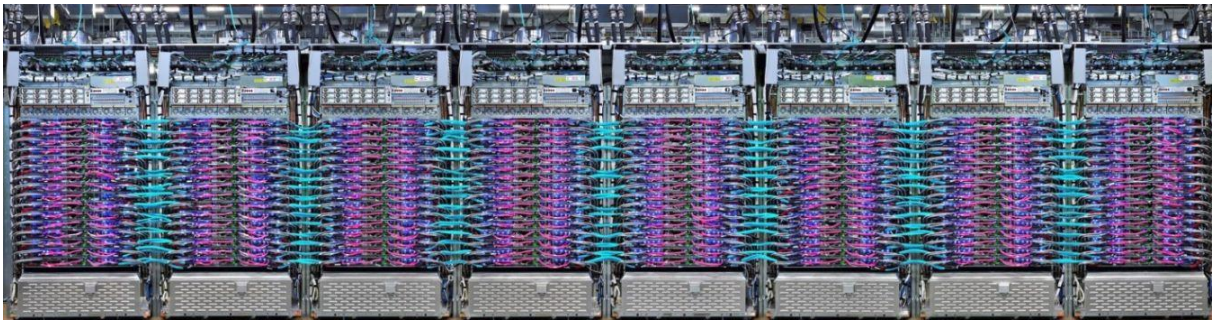

A TPU3.0 pod inside a Google datacentre

Objectives for AI Applications

International efforts at the G7, G20, Organisation for Economic Co-operation and Development (OECD)[6], and the United Nations[7] to build consensus on global AI standards and principles are well developed, supported by concurrent work by international institutions including the Institute of Electrical and Electronics Engineers, and here in Australia by the Office of the Chief Scientist and Standards Australia. We support efforts to align AI for human good, and encourage Australia to participate in global efforts to ensure Australia's voice is heard.

As an organisation working with this technology on a daily basis, Google recognises that the manner in which AI is developed and used will have a significant impact on society, and has thought deeply about how we will approach the use of this technology. We spent more than a year consulting globally with researchers, engineers, governments and nonprofits, seeking input on how this technology should be developed and deployed, and as a result have committed publicly to seven principles that will guide our work. These are not theoretical concepts; they are concrete standards that will actively govern our research, product development, and commercial decisions.

We believe that AI should:

1. Be socially beneficial.

The expanded reach of new technologies increasingly touches society as a whole. Advances in AI will have transformative impacts in a wide range of fields, including healthcare, security, energy, transportation, manufacturing, and entertainment. As we consider potential development and uses of AI technologies, we will take into account a broad range of social and economic factors, and will proceed where we believe that the overall likely benefits substantially exceed the foreseeable risks and downsides.

AI also enhances our ability to understand the meaning of content at scale. We will strive to make high-quality and accurate information readily available using AI, while continuing to respect cultural, social, and legal norms in the countries where we operate. And we will continue to thoughtfully evaluate when to make our technologies available on a non-commercial basis.

2. Avoid creating or reinforcing unfair bias.

AI algorithms and datasets can reflect, reinforce, or reduce unfair biases. We recognize that distinguishing fair from unfair biases is not always simple, and differs across cultures and societies. We will seek to avoid unjust impacts on people, particularly those related to

[6] Beginning at the G7 Information and Communication Ministers Meeting in April 2016, where Japan introduced principles of AI development, leading G7 countries to agree to develop "*AI R&D Principles*" and "*AI R&D Guidelines*," with the cooperation of international organisations such as OECD.

[7] Regionally, through the United Nations Economic and Social Commission for Asia and the Pacific's Committee on Information and Communications Technology & Science, Technology and Innovation in in Bangkok, Thailand; and globally through organisations including the United Nations Interregional Crime and Justice Research Institute's Centre for Artificial Intelligence and Robotics in the Hague.

sensitive characteristics such as race, ethnicity, gender, nationality, income, sexual orientation, ability, and political or religious belief.

3. Be built and tested for safety.

We will continue to develop and apply strong safety and security practices to avoid unintended results that create risks of harm. We will design our AI systems to be appropriately cautious, and seek to develop them in accordance with best practices in AI safety research. In appropriate cases, we will test AI technologies in constrained environments and monitor their operation after deployment.

4. Be accountable to people.

We will design AI systems that provide appropriate opportunities for feedback, relevant explanations, and appeal. Our AI technologies will be subject to appropriate human direction and control.

5. Incorporate privacy design principles.

We will incorporate our privacy principles in the development and use of our AI technologies. We will give opportunity for notice and consent, encourage architectures with privacy safeguards, and provide appropriate transparency and control over the use of data.

6. Uphold high standards of scientific excellence.

Technological innovation is rooted in the scientific method and a commitment to open inquiry, intellectual rigor, integrity, and collaboration. AI tools have the potential to unlock new realms of scientific research and knowledge in critical domains like biology, chemistry, medicine, and environmental sciences. We aspire to high standards of scientific excellence as we work to progress AI development.

We will work with a range of stakeholders to promote thoughtful leadership in this area, drawing on scientifically rigorous and multidisciplinary approaches. And we will responsibly share AI knowledge by publishing educational materials, best practices, and research that enable more people to develop useful AI applications.

7. Be made available for uses that accord with these principles.

Many technologies have multiple uses. We will work to limit potentially harmful or abusive applications. As we develop and deploy AI technologies, we will evaluate likely uses in light of the following factors:

- *Primary purpose and use:* the primary purpose and likely use of a technology and application, including how closely the solution is related to or adaptable to a harmful use
- *Nature and uniqueness:* whether we are making available technology that is unique or more generally available
- *Scale:* whether the use of this technology will have significant impact

- *Nature of Google's involvement:* whether we are providing general-purpose tools, integrating tools for customers, or developing custom solutions

Applications we will not pursue

In addition to the above objectives, Google will not design or deploy AI in the following application areas:

1. Technologies that cause or are likely to cause overall harm. Where there is a material risk of harm, we will proceed only where we believe that the benefits substantially outweigh the risks, and will incorporate appropriate safety constraints.
2. Weapons or other technologies whose principal purpose or implementation is to cause or directly facilitate injury to people.
3. Technologies that gather or use information for surveillance violating internationally accepted norms.
4. Technologies whose purpose contravenes widely accepted principles of international law and human rights.

We found a multi-stakeholder principles-based approach, involving input from an interdisciplinary group of internal and external stakeholders at the design and governance phases, allowed us to understand and address concerns while maintaining technical coherence and preserving space for research innovation that moves us towards implementations that will solve major social, economic and environmental challenges.

Automation and the Australian Workforce

Over the past five years we have seen a plethora of reports released globally touting analyses of jobs that could be lost as a result of the latest wave of technological change, many stemming from a 2013 paper by Frey and Osborne[8] that predicted 47 per cent of U.S. jobs were at risk from automation.

These reports have been of variable quality, and since Frey and Osborne's initial paper the community working on these issues has learned a great deal about how to responsibly investigate this subject. It is worth noting that Osborne's most recent (2017) analysis finds that around one-fifth of workers are in occupations that will likely shrink, with the authors highlighting that the figure is much lower than recent studies of automation have suggested[9].

Motivated by the importance of these issues in long term planning for business leaders and policymakers, and the lack of clarity and Australia-specificity of the existing body of research, Google Australia in 2016 commissioned the economics consultancy AlphaBeta to provide an empirical view of the current state of automation in Australia and its effect on the workforce, drawing on Australian economic statistics.

[8] *The Future of Employment: How susceptible are jobs to computerisation?,* University of Oxford, 2013
[9] *The Future of Skills: Employment in 2030,* Pearson, Nesta, and the University of Oxford, 2017

What we found reflects 'Amara's Law', the adage coined by Stanford's Dr Roy Amara, that states people have a tendency to overestimate the short-term impacts of technological change while underestimating its long-term effects. Our research showed that from 2000 to 2015 the average Australian worker experienced 2 hours of automation across their working week, as routine and repetitive tasks have been automated, and based on current trends will see another 2 hours automated through to 2030[10].

Many of the 80 per cent of Australian workers in service-industry related roles[11] may have to think hard to identify where the productivity gains through automation appeared in their jobs from 2000-2015 - more efficient software for collaboration and service delivery, or communications infrastructure that make coordinating with colleagues in multiple locations more efficient, for example - but at the macro level 2 hours per working week across the economy is hugely significant, and the hours are not distributed evenly, meaning some workers have experienced far more significant change than others.

Differentiating between the workers performing jobs whose task mix is more likely to result in job consolidation and those who are likely to employ the time on a different mix of tasks within their existing role, is essential to policy formation and program delivery.

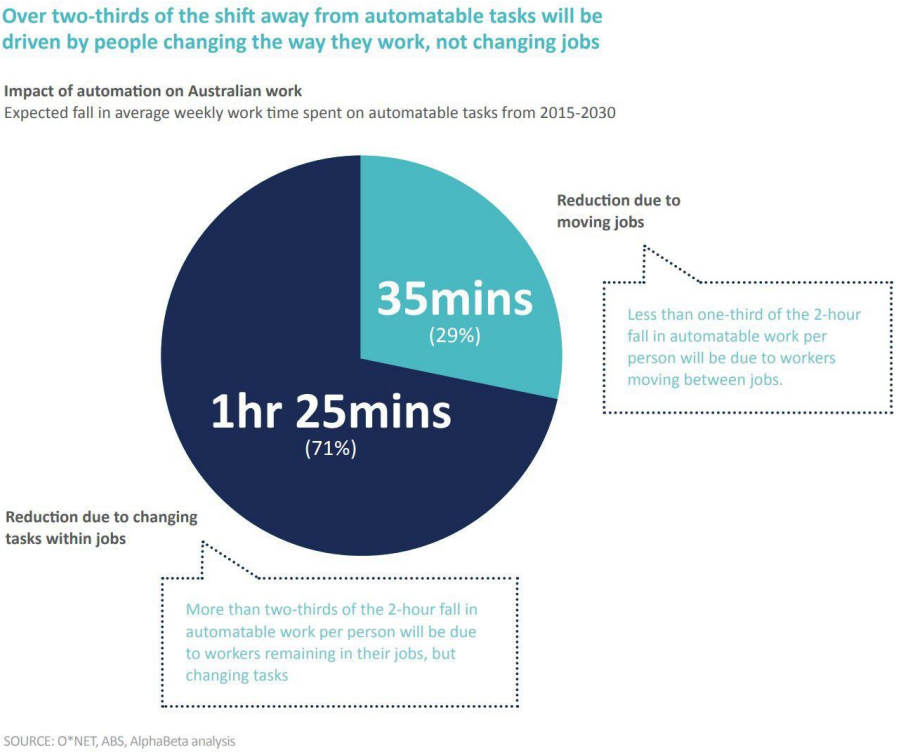

Research shows that routine and predictable tasks, both physical and intellectual, have been automated at higher rates between 2000 and 2015 and we expect are more likely to be automated in the years to 2030. Tasks that have proved more resilient to automation are those that involve interpersonal interaction, decision-making, creativity, and synthesis of information from multiple sources and a degree of qualitative judgement.

[10] *The Automation Advantage*, Alphabeta, 2017
[11] *Australian Industry Report 2016*, Department of Industry, Innovation and Science, 2016

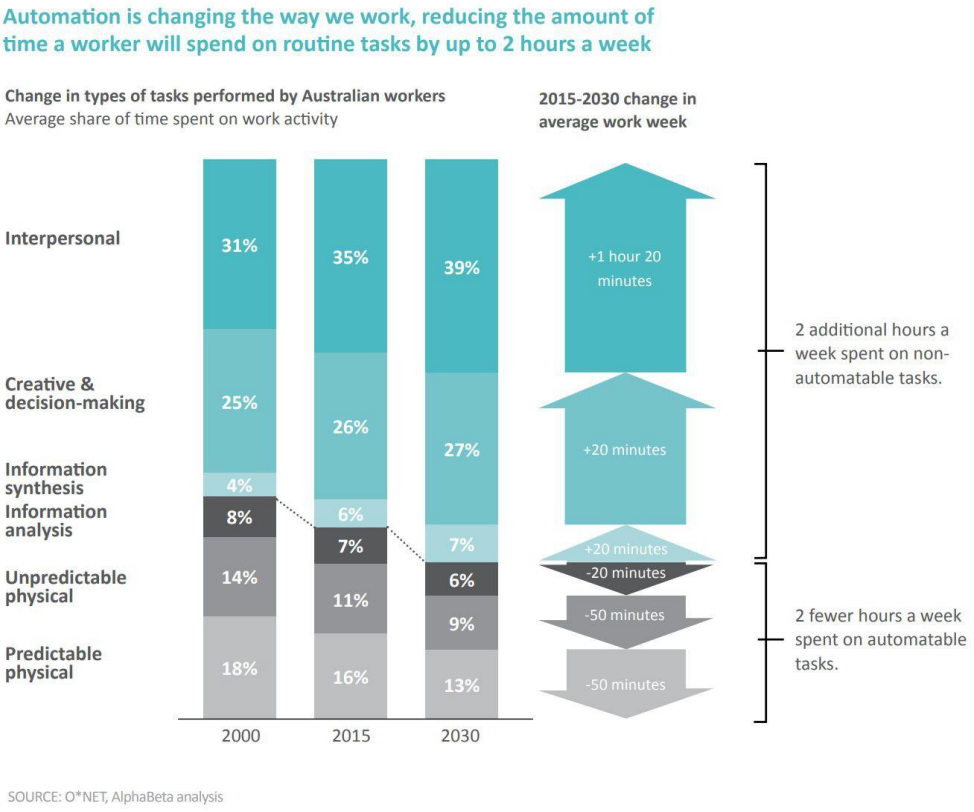

Despite Australia experiencing significant automation across its workforce from 2000-2015, unemployment in Australia remains low[12]. Technology has freed labour capacity and changed the environment in which we live, and new services have appeared to absorb the wealth created by productivity increases, creating new jobs.

If Australian workers in the years from 2015-2030 are able to spend the extra 2 hours of weekly work time that machines are expected to shoulder on higher-value activities (rather than simply reduce their work time by 2 hours per week), it could boost Australia's economy by up to $1.2 trillion in value over that timeframe[13], and different groups of workers affected by technological change have different policy needs.

Highly skilled workers at low-risk of losing their livelihood due to automation are expected to need only minimal government support, and the benefits from automation will likely outweigh its threat. The benefits of accelerating automation and letting these workers naturally shift to higher value work would be substantial for this group, worth $400 billion by 2030.

Workers who perform a large share of automatable tasks may need support to find new ways of working, either in the same jobs or in new ones. An estimated 3.5 million Australian workers are at high-risk of being displaced by automation between 2015 and 2030, and policies providing training and assistance to keep these people in the workforce could yield economic gains worth up to $400 billion.

The costs for society will be highest if Australia fails to adequately prepare its future workers for automation. An additional 6.2 million people are projected to join the Australian workforce

[12] *6202.0 - Labour Force, Australia, Dec 2017*, Australian Bureau of Statistics, 2017
[13] *The Automation Advantage*, Alphabeta, 2017

by 2030, and education policies ensuring these workers are equipped with the right skills could lead to economic gains of $600 billion dollars.

Although the productivity benefits of technology will do the most to raise Australia's economic performance and standard of living over time, we will also see other benefits as a result of the changing nature of work. As automation continues to shift dangerous, tedious and less well remunerated tasks from people to machines, work injuries are set to fall and work satisfaction levels to rise.

- The total number of **workdays lost to injuries sustained from physical work in the Australian economy could fall by 11 per cent** to 1.7 million in 2030.
- Workers currently engaged in more automatable tasks have lower job satisfaction. If current automation trends persist, it is estimated that **62 per cent of low-skilled workers will be happier in their jobs by 2030 compared with today**.
- Australian wage data shows that the least automatable tasks are typically the best paid. **An hour of non-automatable work pays 20 per cent higher wages than an hour of automatable work**.

Our research into the impact of automation on Australia’s economy will be expanded in the coming months of 2018, examining the reskilling challenge and Australia’s stock of education and training. We would welcome the opportunity to brief the Expert Working Group should that be of interest.

ENDS