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 Ross Boyd

Suggested Citation

Boyd, R (2018). Employment and the workforce. Input paper 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.

Input Paper: Horizon Scanning Report on AI for Australian Commonwealth Science Council

Dr Ross Boyd, University of South Australia

This paper will focus on the category of 'Society and the Individual: Employment and the Workforce'. Specifically I will:

- Briefly address a number of issues regarding the relationship of existing and emergent AI-enabled technological developments on productivity and employment/work, these being the focus of much scholarly debate as well as generating considerable public interest and concern;
- Make some observations on: (a) the way contemporary debates over the transformative impacts of AI on work and employment tend to be overwhelmingly, and in many respects unhelpfully, discursively framed in the rather simple (often almost zero-sum) and technologically deterministic terms of whether or not AI/robots are emerging as competitors with, obviators and supplanters of human beings; (b) how this framing rests upon - and indicates the organising power of - a rather misleading and limited conceptualisation of autonomy; and, (c) the outline of an alternative approach which can promise more nuanced and complex understandings of current and prospective transformations, and provide grounds for action in the present and future.

My responses here are based on the analysis published in Boyd and Holton (2017), subsequent elaborations and revisions that have sought to progress this analysis (detailed in papers currently under preparation), and the work of colleagues and myself on a number of research projects, including those receiving ARC and EU funding.

1. Debate over the impacts of AI reprises in many important ways debates over the mechanisation of production that have been ongoing for well over 200 years, in the most recent past in relation to the earlier (1980s, 1990s) introduction of computers into the workplace and the development of the first generation of industrial robots. The leading edge of current debate references emerging AI technical capabilities in areas of machine learning, and flow on gains in, for instance, image recognition, natural language understanding along with prediction techniques and the possibility of advanced machine capabilities in areas of problem solving, self-direction, social interaction, precision dexterity and so forth (Furman and Seaman, 2018; Agrawal, Gans and Goldfarb, 2018; Brynjolfsson, Rock and Syverson, 2017; but also see the contrary view put forward by Marcus, 2018). Equally significant is the integration of physical and digital technologies (the Internet of Things, physical and soft robotics, sensor technologies, nanotechnologies, computer simulation, blockchain, big data, 3D printing) within machine networks (Industry 4.0) that are in many respects self-organising and self-monitoring, highly flexible and, responsive to consumer demand and external/internal disruptions.
2. The contemporary debate turns largely on the question of whether, as a result of these developments, processes of technological innovation and diffusion, and the accompanying transformations to firms, work and employment, will roughly follow the way periodic cycles of technological transformations over the past 200 or so years have played themselves out – meaning the past and the present can serve as reliable guides to the future (Mokyr, Vickers and Diebarth, 2015)? Or, alternatively, does the prospect of unprecedented levels of machine intelligence and autonomy – including the possibility that AI-enabled machines will achieve something resembling independent agency and capabilities once thought uniquely human - along with a rate of change that severely challenges institutional capabilities for

adaptation (Acemoglu and Restrepo, 2016), mean we are on the cusp of experiencing severe, potentially ongoing, disruption to economic organisation, work and employment?

3. One key problem for the current debate is that while the integration of AI-enabled technologies into many aspect of social life is already impressive, many of the technological capabilities identified above are still in the experimental stage, the full range of possibilities they offer (including creative appropriations) and their limits as yet unknown. In those situations where diffusion across workplaces is taking place this is in the early stages, is highly uneven and whatever may (or may not) actually emerge will be contingent upon a range of factors, these irreducible to just technological and/or economic factors. Accordingly, the evidence base for determining trends in productivity, employment and work is underdeveloped and whatever conclusions drawn from it very provisional (Boyd and Holton, 2017).
4. Debate over the productivity impacts of AI/robotics, internationally and in Australia, tends to be dominated by a reprising of the Solow paradox (in 1987 the economist Robert Solow declared that the computer age was evident everywhere save for in the productivity data). As Krishnan, Mischke and Remes (2018) point out the impact of computing on productivity would emerge in the 1990's across a range of economic sectors once ICTs became established as more general purpose enabling technologies: retail and wholesale (through the impacts of computing on supply chain and distribution efficiency), telecommunications and securities trading, as well as computer equipment manufacturing, ICT services and software. With respect of first generation robotics Graetz and Michaels (2015) estimated that, for the 17 countries studied in their sample (using International Federation of Robotics data) robotics contributed around 0.4 to annual GDP growth between 1993 and 2007, this comparing favourably with the introduction of steam engines to the UK. Key sectors here were transport, chemicals and metal working, although they noted a drop off in the pace by the end of the decade they studied as diminishing returns set in. In a recent study Brynjolfsson, Rock and Syverson (2017) argue that the benefits of AI-enhanced technologies will not register in productivity data until, like other general purpose technologies, it prompts the development and implementation of a range of complementary innovations. Krishnan, Mischke and Remes (2018) suggest some of these complementary innovations are currently being diffused, including new digital products and delivery systems. With developments such as the roll out of (semi)autonomous vehicle systems (subject to working through a range of regulatory and technical concerns), or the deployment of smart meters and grids in the utilities sector, they estimate potential productivity growth for the US of around 2% per year over the following decade, with around 60% of that attributable to the new technologies. An alternative account suggests that the current manifestation of the Solow paradox results from the displacement through automation of workers from high productivity/wage manufacturing employment to lower productivity/wage employment in the service sector (Turner, 2018).
5. In a recent review of the (predominantly US) literature on AI and the economy Furman and Seamans (2018) note that AI has "been too small a component of the overall economy to have a significant impact on labour markets". This holds for Australia as well (Hajkowicz et al, 2016). Earlier attempts to estimate the automatability of occupations produced quite alarming figures of close to 50% of US and 38% of UK jobs at risk (Frey and Osbourne, 2013). An Australian replication of this method yielded an estimate of 40% of occupations at risk (Durant-Whyte et al, 2015). Conceptual and methodological shortcomings with the Frey and Osbourne studies have been well documented (see Borland and Coelli, 2017) and need not be rehearsed here. Subsequent analyses have arrived at estimates of 6 – 12 % across 21

OECD countries – with variations attributed to national differences in workplace organisation, investment in automation technologies and education (Arntz, Gregory and Zierahn, 2016) – and around 9% (the OECD average) for Australia (Borland and Coelli, 2017). Nevertheless, the Frey and Osbourne studies did play a significant role in establishing labour displacement as the predominant discursive frame for the debate. Arntz, Gregory and Zierahn also hedge their estimates by suggesting that the implementation of new technologies is a slow process (largely due to legal and societal reasons) and this means: that technological displacement can be inflected in surprising directions; that workers can adjust (learn new skills and transition to new occupations) to changing technological endowments; and, that job losses can be offset by new jobs generated by demand effects (see also Furman and Seamans, 2018). A number of further points can be noted here:

a. While the same cautionary notes in point 3 need to be applied to the rate of change, one thing that needs to be considered is the long term decline in male labour force participation, concentrated among men with no post-secondary education. This trend has been observed in the US (Furman and Seamans, 2018) and Australia (Hajkowicz et al, 2016; 44-45), and suggests that certain groups have struggled to reskill and successfully transition across occupations in response to technological transitions.
b. Hajkowicz et al (2016; 72 – 73) have observed a skills-bias to job generation through technological change, e.g. from 1994 the number of jobs in the traditional printing trades declined by around 17,000, while the numbers of graphic designers (most working with digital technologies) have increased by around 30,000. Elsewhere job polarisation – a hollowing out of middle skill occupations and an increase in non-routine low and high manual and cognitive skill jobs have been found to coincide with increased use of IT and computers for Australia (Borland and Coelli, 2017) and the US (Autor and Salomons, 2018).
c. Automation has also been connected with changing spatial distributions of work, for example in the driverless trains in the Pilbara mines have redistributed work from remote northern western Australia to control centres in Australian capital cities (Ellem, 2016), although conceivably these could be offshored.

6. In recent years there has been a perceptible change in thinking about automation in relation to AI/robotics systems and work, and in particular a shift with respect of the concept (and generative metaphor) of autonomy away from that of fully independent agency towards a new model in which agency is understood as emergent relationship. This has, in part been in response to the emergence of P2P and the gig economy. Ekbia, Nardi and Sabanovic (2015), for example, contrast automated with heteromated systems. Put simply, where automated systems are designed to shift some or many tasks performed by humans over to machines, heteromated systems (Upwork, InnoCentive, Freelancer, Amazons Mechanical Turk and other microwork or crowdsourcing applications, social media, online video games, Paro the robot seal) are designed to work through incorporating end users as indispensable system mediators. While often (ibid) associated with technology-enabled exploitation – the avoidance of modern labour protections and the like – opportunities for the development of a more entrepreneurial model of work are also being explored (Hajkowicz et al, 2016; 37).
7. Other researchers, drawing on their experience working on the development of advanced human-AI/robotics systems, are challenging the autonomy ‘mythology’ they see as inaccurately representing the complex and dynamic human-machine configurations that so-called ‘autonomous’ systems actually entail (Mindell, 2015; Bradshaw, Hoffman, Johnson and Woods, 2013). In doing so these engineers and scientists are in many important respects

making significant moves towards social scientific – especially social theoretical, sociological and anthropological – understandings of autonomy and agency. Drawing on Susan Leigh Star's work on infrastructures, Lucy Suchman (2007) has, for example, demonstrated how the autonomy of robotic agents is an effect of a complex surrounding support structure – including the performances of other human agents. More recently field research on US Airforce Drone operations by M.C. Elish (2018) at the Data and Society Research Institute, New York, "suggests that human infrastructures of unmanned drone operations are obscured while also remaining quintessential to their operation. As drones promise a path towards autonomous AI systems, such autonomy only emerges by masking the human labor and networks that create and maintain it" (p2).

8. The conceptualisation of autonomy as an emergent relation offers considerable scope for shifting the discursive frame of discussion and debate away from a 'displacement' model that is of limited accuracy or usefulness. In a recent analysis Boyd and Holton (2017) stressed the importance of critically evaluating the range of discourses about technological change, "because stories about technological change do not simply sit above technological innovation and diffusion but help to constitute and direct or re-direct change itself." This is because they serve as organising visions for not just the public imagination but also researchers in the field (Natale and Ballatore, 2017) along with those designing public policies and corporate strategies. And shifting the discursive frame of debate in this way enables machine learning to be, as Stilgoe (2018) suggests, more closely aligned with ongoing processes of social learning about emergent technologies and their actual and potential impacts.

References:

Acemoglu, Daron, and Pascual Restrepo. 2016. "The race between machine and man: Implications of technology for growth, factor shares and employment." National Bureau of Economic Research working paper No. w22252.

Agrawal, Ajay, Joshua S. Gans and Avi Goldfarb. 2018. Prediction Machines: The Simple Economics of Artificial Intelligence. Cambridge, MA: Harvard Business Review Press.

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, and anna salomons. 2018. "is automation labor displacing? productivity growth, employment and the labor share" brookings papers on economic activity

Borland, Jeff and Michael Coelli. 2017. "Are Robots Taking Our Jobs?" Australian Economic Review, 50(4), 377-97

Boyd, Ross and Robert J. Holton. 2017. "Technology, Employment, Innovation, and Power: Does Robotics and Artificial Intelligence Really Mean Social Transformation", Journal of Sociology, doi.org/10.1177/1440783317726591

Bradshaw, Jeffrey M., Robert R. Hoffman, Mathew Johnson and David D. Woods. 2013. "The Seven deadly Myths of 'Autonomous Systems'", IEEE Intelligent Systems, May/June, 54-61.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2017. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In Ajay Agrawal, Joshua S. Gans and Avi Goldfarb (eds). The Economics of Artificial Intelligence. Chicago: University of Chicago Press

Durant-Whyte, Hugh et al. 2015. "The impact of computerisation and automation on future employment. Australia's Future Workforce. Melbourne: CEDA

Ekbia, Hamid, Bonnie Nardi and Selma Sabanovic. 2015 "On the Margins of the Machine: Heteromation and Robotics". In IConference 2015 Proceedings.

Ellem, Bradon. 2016. "Geographies of the labour process: automation and the spatiality of mining." Work, Employment and Society, 30(6). 932-948.

Elish, M.C. 2018. "(Dis)Placed Workers: A Study in the Disruptive Potential of Robotics and AI". WeRobot 2018 Conference

Frey, Carl B. and Michael A. Osborne. 2013. "The Future of Employment: How Susceptible are Jobs to Computerisation", at www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf

Furman, Jason and Robert Seamans. 2018. "AI and the Economy." National Bureau of Economic research Working Paper 24689

Graetz, Georg and Guy Michaels. 2015. "Robots at Work." Centre for Economic Performance Discussion Paper No. 1335

Hajkowicz, Stefan et al. 2016. Tomorrow's Digitally Enabled Workforce. Brisbane: CSIRO

Krishnan, Mekala, Jan Mischke, and Jaana Remes. 2018. "Is the Solow Paradox Back?" McKinsey Quarterly, June 2018. https://www.mckinsey.com/business-functions/digital-mckinsey/our-insights/is-the-solow-paradox-back

Marcus, Gary. 2018. "Deep Learning: A Critical Appraisal." arXiv:1801.00631

Mindell, David. 2015. Our Robots, Our Selves: Robotics and the Myths of Autonomy. New York: Viking

Mokyr, Joel, Chris Vickers and Nicholas L. Diebarth. 2015. "The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different", Journal of Economic Perspectives, 29(3), 31-50

Natale, Simone, and Andrea Ballatore. 2017. "Imagining the thinking machine: Technological myths and the rise of artificial intelligence. Convergence. 1-16.

Stilgoe, Jack. 22018. "Machine learning, social learning and the governance of self-driving cars. Social Studies of Science. 48(1). 25-56.

Suchman, Lucy. 2007. Human-Machine Reconfigurations. Cambridge: Cambridge University Press

Turner, Adair. 2018 'Capitalism in an Age of Robots", Centre for New Economic Thinking, https://ineteconomics.org/uploads/papers/Papers-Turner-Capitalism-In-The-Age-Of-Robots.pdf