Horizon Scanning Series

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

Training the Next Generation of AI Researchers

This input paper was prepared by Mark Reynolds

Suggested Citation

Reynolds, M (2018). Training the Next Generation of AI Researchers. 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.

Training the next generation of AI researchers

Mark Reynolds mark.reynolds@uwa.edu.au

The University of Western Australia Perth, Australia

Abstract. How will the training of the next generation of AI researchers be different? We consider this from an Australasian perspective.

Introduction

The Australian Council of Learned Academies have been commissioned by Australia's Chief Scientist and the Commonwealth Science Council to undertake a comprehensive study into the opportunities and challenges presented by deployment of artificial intelligence in Australia and New Zealand. In collaboration with the Royal Society of Te Ap¯arangi, the study, through its final report, is intended to inform policy considerations into the social, legal, ethical, technological and economic implications for broader use of artificial intelligence applications.
I have been asked to put together a short paper for this report about the challenges in training future AI researchers. Specifically, I have been asked to examine the following:

Do we need to train more researchers to develop fundamental science of AI, or to apply AI to new areas?
How will AI affect the process of research across other disciplines?
How will AI impact the analysis of big data?
How does the training of AI researchers need to change?

Do we need to train more researchers to develop fundamental science of AI, or to apply AI to new areas?

There are shortages of workers in all areas of IT. [9] "Demand for ICT workers is set to grow by almost 100,000 to 758,700 workers by 2023" but "with fewer than 5,000 domestic ICT graduates a year, the only way we'll reach workforce targets is by importing labour, much as we've done for the past five years. We need more ICT workers with skills in artificial intelligence, data science, cyber security and blockchain, and filling these positions with migrants suggests a missed opportunity to provide rewarding employment for the next generation of Australian workers."

The fundamental science of AI includes such topics as computer vision, natural language processing, creativity, optimisation, search, machine learning, transparency, reasoning, planning, common sense, strategy and robotics. There are still huge challenges in these foundations. Computer vision and image processing, for example, is a major research topic with vast conferences filled with papers tackling major concerns. CVPR (Conference on Computer Vision and Pattern Recognition) 2018, the premier annual computer vision event, held in June 2018 in Salt Lake City, had nearly 1000 papers ranging from hand and pose tracking to scene reconstruction, and advances in all the machine learning techniques that support such tasks [14]. And there is a long list of other similarly vast high quality international annual conferences in the vision area. Other foundational areas of AI are possibly not quite as hyperactive but they all are attracting serious research effort at the moment.

And how long will this research be needed? According to David L. Waltz, Vice President of Computer Science Research at the NEC Research Institute, “The pursuit of the ultimate goals of AI continues to have practical consequences in the form of new industries, enhanced functionality for existing systems, increased productivity in general, and improvements in the quality of life. But the ultimate promises of AI are still decades away, and the necessary advances in knowledge and technology will require a sustained fundamental research effort” [33].

Commercial spending on AI systems is exploding expected to be up from $12.5 billion in 2017 to more than $46 billion by 2020 [19]. ”Double-digit spending growth is expected for cognitive and artificial intelligence systems across all industries”. Currently the big users are banks and securities tackling risk and fraud identification but also to better match products with clients. Big growth in the next few years will come from a wide range of other industries from manufacturing and retail to healthcare. To support the development of the applied AI technology across industry, government and society in general there will be a huge demand for researchers in AI applications.

In [34], the Brookings institute lists novel AI applications in finance, national security, health care, criminal justice, transportation, and smart cities. Development of these applications all need domain expertise as well as AI expertise. Development of such applications is already well underway in some industries such as mining [32] but many more applications are still some way in the future— AI in the education industry [31], for example.

The increased use of AI in society, and new uses of AI, also raises important ethical questions. These will need answering. Researchers in IT and a variety of other disciplines will be engaged to tackle these. According to Julia Bossmann, writing for the World Economic Forum, “In many ways, this is just as much a new frontier for ethics and risk assessment as it is for emerging technology” [3]. Major issues include how we: employ humans when jobs are replaced; distribute the wealth created by machines; are affected by machines; can guard against mistakes; eliminate AI bias; keep AI safe from adversaries; protect against unintended consequences; stay in control of a complex intelligent system; and define the humane treatment of AI.

It is worth noting here that “transparency”, the explainability of AI decisions, particularly as they affect people’s lives, is both an AI research question for the technical and underlying fundamental science of AI, and also a major ethical concern [5]. AI technology developers realise there are many advantages to being able to understand and explain the individual decisions of a trained machine learning algorithm, but many of the algorithms that are the best at making predictions are essentially “black boxes” [16]. Working towards transparency is a very active research interest in AI [30].

How will AI affect the process of research across other disciplines?

Artificial Intelligence will drive big changes in the way that academic and industrial research is done just as it will significantly affect many human activities. There will be (1) discipline specific changes, there will also be (2) similar sorts of changes across disciplines which have similar concerns and finally there will be
(3) very general changes to the basic processes of research that will be noticed across a wide range of disciplines.

In the category of discipline specific changes we have machine learning doing repetitive, but sometimes subtle, identification tasks such as (in physics) telling the difference between a shower of photons and a shower of electrons [22]; repetitive and meticulous tasks such as smoothing out noisy astronomical pictures of galaxies using a generative adversarial network [22]; complex and time consuming searching such as in chemistry finding the right sequence of reactions to create small organic molecules, such as drug compounds [28]; or following a group of complex and interacting rules such as in genetics, predicting how neighbouring genes affect the behaviour of genes that then cause autism [22].

In this category we may also find examples where certain AI techniques open up whole new research fields within a discipline. Research into using machine learning, or advanced optimisation or search techniques for the development of the very data intensive field of “precision medicine” would fit in here [17]. So AI will change the process of medical research because there are a whole new set of concerns, techniques, aims, etc. What is researched will change, so the process of research will need corresponding changes. AI in education [24] is another huge sub-field opening up.

The second category of affects are the more general affects will be in those sets of disciplines which together make use of certain techniques or together have similar concerns. Let us list a few examples.

Image processing algorithms have a fantastic range of use cases, but just the relatively simple task of counting appearances of similar objects allows the automation of what was a very labor intensive part of typical research processes in many disciplines. Now, to choose just one example from many, we do not need to employ graduate biology students to spend hours counting fish in videos [26]. Natural language processing such as analysis of text, or language can be applied to sift through maintenance records, or transcripts of conversation or meticulous records of human observations. This will again allow the reassignment of many hours of researcher labour but also the open up new avenues of research in many areas from education and psychology to system health engineering. At the University of Pennsylvania’s Positive Psychology Centre, for example, psychologists, doctors, and computer scientists use machine learning and natural language processing on social media data to detect patterns in the emotional and physical health of the population [22].

Another branch of AI, artificial reasoning, will increasingly be used in math-semantics, philosophy [2], science, and law [21] to go beyond human capabilities in stamina of rigorous deduction.

Artificial creativity is a longer term goal but will be used in industrial design, architecture, engineering and art. Creativity is needed by practitioners here but also therefore by researchers trying to improve or understand the practice. There are (justifiable) claims that AI is not very good at the creativity at the moment (see for example [24]) but there is significant research work on this and already AI techniques are quite capable of coming up with novel approaches: “for example, genetic algorithms explore large design spaces and discover solutions human engineers would not imagine (or like)” [11].

The third category of AI affect on research practice are the most general sorts of affects which will be caused by AI technologies which address some fundamental part of research such as the search for relevant literature. AI research assistants such as Iris [20] and IBM Watson [18] can already look through the vast literature of previous research studies or depths of organisations’ knowledge for relevant facts.

Another fundamental part of doing research, particularly in the Sciences, is the act of formulating a hypothesis. To some extent this is a creative process, but it is also a search through the realm of potential hypotheses looking for one that describes the observations before the researcher. AI is already at work doing this in limited ways building on its abilities in data mining, search and pattern recognition [27].

Research publishing itself will also be affected. “Software with the capability to complete subject-oriented reviews of manuscripts, coupled with automated methodological reviews, could enable a fully automated publishing process – including the decision to publish” [10].
We may even eventually see AI writing research papers [12].

How will AI impact the analysis of “big data”?

Over the last decade the term “Big Data” has come to be used for applications and situations in which traditional data management techniques are challenged by the 3 Vs, the Volume, Variety and Velocity of the data arriving to be able to be used for insight and decision making [15]. Although, this is far from being an unambiguous or universally accepted definition, it is clear that huge and increasing amounts of various types of data are available to be collected and made use of, are being collected, are having to be collected (for audit purposes) and are able to be processed usefully. And the uses range across the whole spectrum of human activity.
AI does, and will increasingly, help with analysis of big data.

Although specific machine learning and data mining techniques are not al- ways scalable to cope with huge amounts of data, development of more scalable alternatives is a very active part of AI research. For example, the classical method of doing linear regression (the Normal Equation) is often too slow to be practical when there are a large number of data features but there are alternative techniques to use in such circumstances (such as Stochastic Gradient Descent) [16]. The computational complexity and relative merits of various rival techniques is a well understood factor by data engineers when they make their design decisions. And AI researchers are constantly motivated by Big Data applications to find faster ways to do machine learning. The algorithms, better hardware and network solutions are deployed by software such as TensorFlow [16] which was designed to allow convenient implementation of large scale machine learning. Such work is current and ongoing. Applications are widespread including search- ing for valuable and interesting patterns in large business customer databases, industrial operations records and data intensive science.

A huge amount of new digital data, and an ever increasing amount, is being collected in the form of recorded video. Much more is being viewed by cam- eras and being streamed without being recorded in any long term sense. For example, webcams and CCTV target freeways, city squares, bridges and cycle routes, security monitors workplaces, car parks, and shopping centres, and deci- sion makers watch video feeds from industrial processes. AI work such as Human Action Recognition [1] helps allow automated video annotation. This in turn, allows automated alerts for suspicious behaviour, dangerous situations, anoma- lies, classification for compiling statistics and analyses, and subsequent search for tracking, forensics, etc. The technology here is developing rapidly.

Large amounts of textual data is also being collected or is available to be col- lected easily in the future and AI natural language processing and text data min- ing techniques are being applied. Data comes from commercial databases which often include text commentary, automatically generated transcripts of recorded audio from call centres, customers and workplaces, maintenance records, company reports, wikipedia, news reports, and of course social networks, Techniques applied include text analytics, information extraction, text summarisation, question answering and sentiment analysis [15].

Combining the natural language processing, sentiment analysis with AI net- work analysis gives us Social Media analytics. Already we have companies such as IBM offers “social merchandising” for retail and consumer products using Twitter data [25].

In a separate application, machine learning and AI optimisation techniques can also be used to manage the workflow of large data processing tasks such as the massive volume of constantly delivered radio astronomy data to be arriving soon via the Square Kilometre Array (SKA) in Western Australia [35].

How does the training of AI researchers need to change?
We have already seen above that one desperately needed change to AI education is that there needs to be more of it: we need to involve more students, attract more capable students and graduate more AI practitioners, researchers and experts. However, the vast majority of AI development and research needs to come from students with relevant post-secondary training and the majority of such students are going to need good secondary STEM backgrounds, with strong mathematical abilities to the forefront. It is not news that Australia has a problem in this pipeline [23].

But how else does AI training need to change?

Bostrom [4], for example, makes a list of the sorts of requirements that will be placed on future AI technology by society for ethical reasons. They include the need for the AI to be: transparent to inspection, predictable to those they govern, robust against manipulation, and responsible (deployed in a context in which we know who takes responsibility). AI researchers will need to understand these “constraints”. As [29] mentions, we are already seeing that autonomous vehicles, as just one example of AI, bring their own truckload of ethical issues, along these dimensions. AI training needs to incorporate ethics.

As AI spreads throughout society, AI researchers need to learn how to work with other discipline experts to investigate, develop, design and build systems. Sometimes other disciplines can be seen as clients of the AI developer, for ex- ample, perhaps an astrophysicist wanting to use machine learning to identify galaxies from amongst a huge collection of images. However, even there, the AI researcher, like any software developer or applied statistician needs some do- main knowledge, and to know how to speak with the domain expert. But also the AI systems will need to work with humans, work with more standard soft- ware, work within organisations, work within society. The AI system developer needs to know at least a little about and know how to work with experts from fields such as the following: software engineering, human-computer interaction (or interaction design), information systems, business, psychology, economics, politics, industrial relations, human resource management, and ethics. For all these reasons AI education needs to become more interdisciplinary.

Including diversity of thought in developing AI is essential for society to utilize it successfully [24]. Rexford reminds us that “there’s certainly a significant underrepresentation problem in all areas of STEM [science, technology, engineering, and math], and this is higher in computer science than in some other areas and higher in artificial intelligence than in general computer science. So diversity is definitely a significant problem. And there’s a tremendous shortage of skilled people to do this work. There’s this general cry of ‘We need more people, so we really can't afford to leave talent on the table'." But Rexford goes on "there's a second piece to this: the technology being created is so important in so many lines of work and in so many businesses that leaving a large part of the population out of the conversation leaves these people ineffective in a whole bunch of areas in our society — not just computer science, but everywhere." There are many examples of how new technology for use across the population goes wrong when the design team lacks diversity [8].

"Designers and researchers from a range of disciplines need to conduct what we call social-systems analyses of AI, need to assess the impact of technologies on their social, cultural and political settings" [8]. Looming large here are the consequences for (human) employment opportunities resulting from large scale deployment of AI across industry [29]. AI researchers need to be aware of potential social impacts.

Understanding legal compliance is another need. For example, the EU's new General Data Protection Regulation (GDPR) legislation [13] article 15 grants the data subject the right to meaningful information about the logic involved and the envisaged consequences of such processing for the data subject when automated decision-making is used [7]. So explainability, justifiability and transparency come up again.

To end on some good news, AI work in improving accessibility of the digital world [6] and in supporting more engaging educational experiences [24] is avail- able to help in leading more young people into studying and working towards a career in AI research. Let us make the most of it.

References

1. J. Aggarwal and M. Ryoo. Human activity analysis: A review. ACM Comput. Surv., 43(3):16:1–16:43, 2011.
2. Christoph Benzmu¨ller and Bruno Woltzenlogel Paleo. Automating G¨odel's onto-logical proof of god's existence with higher-order automated theorem provers. In Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan, editors, ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014), volume 263 of Frontiers in Artificial Intelligence and Applications, pages 93–98. IOS Press, 2014.
3. Julia Bossmann. Top 9 ethical issues in artificial intelligence. https://www.weforum.org/agenda/2016/10/top-10-ethical-issues-in-artificial- intelligence/. Accessed: 2018-07-21.
4. N. Bostrom and E. Yudkowsky. The ethics of artificial intelligence. In The Cambridge Handbook of Artificial Intelligence, pages 316–334. CUP, 2014.
5. Davide Castelvecchi. Can we open the black box of ai? https://www.nature.com/news/can-we-open-the-black-box-of-ai-1.20731. Accessed: 2018-07-21.
6. Joe Chidzik. 5 ways ai could transform digital accessibility. https://www.abilitynet.org.uk/news-blogs/5-ways-ai-could-transform-digital- accessibility. Accessed: 2018-07-21.
7. Intersoft Consulting. The EU General Data Protection Regulation, article 15. https://gdpr-info.eu/art-15-gdpr/. Accessed: 2018-07-21.
8. Kate Crawford and Ryan Calo. There is a blind spot in AI research. https://www.nature.com/news/there-is-a-blind-spot-in-ai-research-1.20805. Accessed: 2018-07-21.
9. Deloittes. ACS Australia's digital pulse. https://www.acs.org.au/content/dam/ acs/acs-publications/aadp2018.pdf. Accessed: 2018-07-21.Chadwick C. DeVoss.

Artificial intelligence can expedite scientific communication and eradicate bias from the publishing process.
http://blogs.lse.ac.uk/impactofsocialsciences/2017/05/11/artificial-intelligence- can-expedite-scientific-communication-and-eradicate-bias-from-the-publishing- process/. Accessed: 2018-07-21.
10. Allen B. Downey. Think Complexity, volume 2. Green Tea Press, 2016.
11. Times Higher Education. Rise of the research-bots: Ai software that writes your papers for you. https://www.timeshighereducation.com/news/rise-research-bots- ai-software-writes-your-papers-you. Accessed: 2018-07-21.
12. EUGPDR.org. The EU General Data Protection Regulation.
13. https://www.eugdpr.org. Accessed: 2018-07-21.
14. Computer Vision Foundation. CVPR 2018 open access repository. http://openaccess.thecvf.com/CVPR2018.py. Accessed: 2018-07-21.
A. Gandomi and M. Haider. Beyond the hype: Big data concepts, methods, and
15. analytics. International Journal of Information Management, 35(2):137–144, 2015.
16. Aurelien Geron. Hands-On Machine Learning with Scikit-Learn and TensorFlow. O'Reilly, 2017.
17. Richard Hodson. Precision medicine. Nature, 537(S49), 2016.
18. IBM. Put AI to work. https://www.ibm.com/watson/. Accessed: 2018-07-21.
19. International Data Corporation (IDC). Worldwide spending on cognitive and arti- ficial intelligence systems forecast to reach $12.5 billion this year, according to new idc spending guide. https://www.idc.com/getdoc.jsp?containerId=prUS42439617. Accessed: 2018-07-21.
20. Iris.ai. The future of science. https://iris.ai. Accessed: 2018-07-21.
21. Bernard Marr. How ai and machine learning are transforming law firms and the legal sector. https://www.forbes.com/sites/bernardmarr/2018/05/23/how-ai-and- machine-learning-are-transforming-law-firms-and-the-legal-sector. Accessed: 2018- 07-21.
22. Science News. AI is changing how we do science. Get a glimpse.
23. http://www.sciencemag.org/news/2017/07/ai-changing-how-we-do-science- get-glimpse. Accessed: 2018-07-21.
24. Office of the Chief Scientist. Australia?s stem workforce: a survey of
25. employers. https://www.chiefscientist.gov.au/wp-content/uploads/DAE OCS-Australias-STEM-Workforce FINAL-REPORT.pdf. Accessed: 2018-07-21.
26. Jennifer Rexford and Rik Kirkland. The role of education in AI (and vice
27. versa). https://www.mckinsey.com/featured-insights/artificial-intelligence/the- role-of-education-in-ai-and-vice-versa. Accessed: 2018-07-21.
28. J. Ryoo and E.J. Kim. Big data e-book. http://sites.psu.edu/bigdataebook/.
29. Accessed: 2018-07-21.
A. Salman, A. Jalal, Faisal Shafait, Ajmal Mian, M. Shortis, J. Seager, and E. Har- vey. Fish species classification in unconstrained underwater environments based on deep learning. Limnology and Oceanography : Methods, 14(9):570–585, 9 2016.
30. Valentina Sasselli and Hylke Koers. How big data and AI can help you generate
31. your scientific hypothesis. https://www.elsevier.com/connect/how-big-data-and- ai-can-generate-your-scientific-hypothesis. Accessed: 2018-07-21.
32. M. H. S. Segler, M. Preuss, and M. P. Waller. Planning chemical syntheses with
33. deep neural networks and symbolic AI. Nature, 555:604–610, 2018.
34. Alex Shashkevich. Stanford scholars, researchers discuss key ethical ques- tions self-driving cars present. https://news.stanford.edu/2017/05/22/stanford- scholars-researchers-discuss-key-ethical-questions-self-driving-cars-present/. Ac- cessed: 2018-

07-21.
35. Sarah Tan, Rich Caruana, Giles Hooker, and Yin Lou. Auditing black-box models using transparent model distillation with side information. CoRR, abs/1710.06169v3, 2018.
36. Karl Utermohlen. 4 ways ai is changing the education industry. https://towardsdatascience.com/4-ways-ai-is-changing-the-education-industry-b473c5d2c706. Accessed: 2018-07-21.
37. Jon Walker.Ai in mining: Mineral exploration, autonomous drills, and more. https://www.techemergence.com/ai-in-mining-mineral-exploration- autonomous-drills/. Accessed: 2018-07-21.
38. David L. Waltz. Artificial intelligence: Realizing the ultimate promises of comput- ing. https://homes.cs.washington.edu/ lazowska/cra/ai.html. Accessed: 2018-07- 21.
39. Darrell M. West and John R. Allen. How artificial intelligence is transform- ing the world. https://www.brookings.edu/research/how-artificial-intelligence-is- transforming-the-world/. Accessed: 2018-07-21.
40. Chen Wu, Andreas Wicenec, and Rodrigo Tobar. Partitioning SKA dataflows for optimal graph execution. CoRR, abs/1805.07568, 2018.