Expert and Non-expert Opinion About Technological Unemployment

Toby Walsh[1,2,3]

[1] University of New South Wales, Sydney, Australia

[2] Data61, Locked Bag 6016, UNSW, Kensington, Sydney, Australia

[3] Technical University of Berlin, Berlin, Germany

Abstract: There is significant concern that technological advances, especially in robotics and artificial intelligence (AI), could lead to high levels of unemployment in the coming decades. Studies have estimated that around half of all current jobs are at risk of automation. To look into this issue in more depth, we surveyed experts in robotics and AI about the risk, and compared their views with those of non-experts. Whilst the experts predicted a significant number of occupations were at risk of automation in the next two decades, they were more cautious than people outside the field in predicting occupations at risk. Their predictions were consistent with their estimates for when computers might be expected to reach human level performance across a wide range of skills. These estimates were typically decades later than those of the non-experts. Technological barriers may therefore provide society with more time to prepare for an automated future than the public fear. In addition, public expectations may need to be dampened about the speed of progress to be expected in robotics and AI.

Keywords: Survey, technological unemployment, artificial intelligence (AI).

1 Introduction

Areas like deep learning are advancing artificial intelligence rapidly[1]. The World Economic Forum has predicted that we are at the beginning of a Fourth Industrial Revolution which will transform the nature of our economies and eliminate many current occupations[2]. At the same time, new technologies will also create many new occupations. It remains an open question whether more jobs will be created than destroyed. Back in 1930, Keynes[3] predicted that technological changes of the Second Industrial Revolution would eventually create more jobs. He was correct as unemployment rates are now lower than they were before. However, this may not be the case in the future as we are likely to have fewer and fewer advantages over the machines.

In any case, it is likely that the new occupations created will require different skills to those destroyed. For instance, autonomous vehicles will probably be commonplace on our roads within the next few decades. Taxi and truck drivers will therefore need other skills than just the ability to drive if they are to remain employed. It is thus an important question for our societies in preparing for this future of technological change to understand the occupations at risk of automation.

2 Background

In 2013, a much reported study by Frey and Osborne[4] estimated that 47% of total employment in the United States was under risk of automation in the next two decades. Ironically, the study used machine learning to predict occupations at risk. Even the occupation of predicting occupations at risk from automation has been partially automated.

Subsequent studies have reached similar conclusions. For instance, similar analysis has estimated that 40% of total employment in Australia is at risk of automation[5], and even larger figures for developing countries like China at 77% and India at 69%[6].

Frey and Osborne suggested three barriers to automation: Occupations requiring complex perception or manipulation skills, occupations requiring creativity, and occupations requiring social intelligence. Computers are significantly challenged in these three areas at present and may remain so for some time to come.

Frey and Osborne's study used a training set of 70 occupations from the O*Net database of U.S. occupations. This training set was hand labelled by a small group of economists and machine learning researchers at a workshop held in the Oxford University Engineering Sciences Department. Classification was binary. Each occupation was classified either at risk in the next two decades from automation or not. Labels were only assigned to occupations where there was confidence in the classification.

We do not wish to discuss here whether the O*Net database provides features adequate to extrapolate to the full set of 702 occupations. This is a difficult question to address as we do not have a gold standard of occupations actually at risk. Their classifier did, however, perform well on the training set with a precision (positive predictive value) for occupations at risk of automation of 94%, a sensitivity of 81%, and a specificity of 94%. We also leave it as future work to extrapolate from jobs at risk to percentage of workforce unemployed.

We focus here on the training set of 70 occupations used in [4]. This study hand labelled 37 of these 70 occupations as being at risk of automation (53%). The final accuracy of the classification of 702 occupations depends critically on the accuracy with which this smaller training set was hand labelled.

This training set was chosen as it could be classified "with confidence". We therefore gave this training set to three much larger groups to classify: experts in AI, experts in robotics and, as a comparison, non-experts interested in the future of AI. In total over, we sampled over 300 experts and 500 non-experts. Our survey is the largest of its kind ever performed.

3 High level machine intelligence

In addition to classifying the training set, we asked both the experts and the non-experts to estimate when computers might be expected to achieve a high-level of machine intelligence (HLMI). This was defined to be when a computer might be able to carry out most human professions at least as well as a typical human. In 2012/2013, Müller and Bostrom[7] surveyed 170 people working in AI to predict when HLMI might be achieved.

As there is significant uncertainty as to when HLMI might be achieved, they asked when the probability of HLMI would be 10%, 50% and 90%. The median response for a 10% probability of HLMI was 2 022, for a 50% probability was 2 040, and for a 90% probability was 2 075. We wanted to see if people who were more cautious at predicting when HLMI was likely to be achieved were also more cautious at predicting occupations at risk of automation.

We also wished to update and enlarge upon Müller and Bostrom's survey. Given some of the high profile advances made recently in subareas of AI like deep learning[8], it might be expected that HLMI would be predicted sooner now than back in 2012/2013. We also wanted to survey a much larger sample of experts in AI and robotics than Müller and Bostrom.

Only 29 of the 170 who answered Müller and Bostrom's survey were leading experts in AI, specifically 29 members of the 100 must cited authors in AI as ranked by Microsoft Academic Research. The largest group in their survey were 72 participants of a conference in artificial general intelligence (AGI). This is a specialized area in AI where researchers are focused on the question of building general intelligence. Much research in AI is, by comparison, focused on programming computers to do very specialized tasks like playing Go[9] or interpreting mammograms[10] and not on building general purpose intelligence.

Researchers in AGI might be expected to be pre-disposed to the early arrival of HLMI. Indeed the AGI group were the most enthusiastic to complete Müller and Bostrom's survey. 64% of the delegates from this AGI conference completed the survey, compared to an overall response rate of just 31%. In addition, the AGI group typically predicted HMLI would arrive earlier than the other respondents to the survey. We conjectured that experts in AI and robotics not focused on AGI would be more cautious in their predictions.

More recently in March 2016, Oren Etzioni[11] wanted to test a similar hypothesis about Müller and Bostrom's results. To do so, he sent out a survey to 193 Fellows of the Association for the Advancement of Artificial Intelligence (AAAI). In total, 80 Fellows responded (41% response rate). Respondents included many leading researchers in the field like Geoff Hinton, Ed Feigenbaum, Rodney Brooks, and Peter Norvig.

Unfortunately, Etzioni's survey asked a different and simpler question ("When do you think we will achieve Superintelligence?" where Superintelligence is defined to be "an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom and social skills"). Etzioni's survey also only offered 4 answers to the question of when Superintelligence would be achieved (in next 10 years, 10–25 years, more than 25 years, never).

It is difficult to compare the results of Etzioni's survey with Müller and Bostrom's. None of the AAAI Fellows responding selected "in the next 10 years", 7.5% selected "in the next 10–25 years", 67.5% selected "in more than 25 years", and the remaining 25% selected "never". If Etzioni's question is equated with Müller and Bostrom's question about a 90% probability of HLMI, then the responses of the two surveys appear to be similar. However, it is very difficult to draw many conclusions given the rather ambiguous question, and the larger granularity on the answers.

4 Methods

Our survey was performed between the 20th January, 2017 and the 5th February, 2017. The survey involved three distinct groups. The first group were authors from two leading AI conferences: the *Annual conference of the Association for the Advancement of Artificial Intelligence* (AAAI 2015), and the *International Joint Conference on Artificial Intelligence* (IJCAI 2011). Both conferences are highly selective and publish some of best new work in AI. 200 authors from this group completed our

survey.

The second group consisted of IEEE Fellows in the IEEE Robotics & Automation Society and authors of a leading robotics conference: *IEEE International Conference on Robotics and Automation* (ICRA 2016). This is also a highly selective conference that publishes some of the best work in robotics. In total, 101 people from this group completed the survey.

The third and final groups surveyed were readers of an article from the website "The Conversation". This Australian and British website publishes news stories and expert opinion from the university sector, and is partnered with Reuters and the Press Association. The article containing the link to the survey was entitled "Know when to fold 'em: AI beats world's top poker players".

The article discussed the recent victory of the CMU Libratus poker program against some top human players. It used this as an introduction to the Frey and Osborne report on tasks that could be automated. It ended by inviting readers to help determine the "wisdom of the crowd" by completing the survey. There were 548 responses in this third group.

The readers of The Conversation have the following geographical distribution: 36% Australia, 29% United States, 7% United Kingdom, 4% Canada, and 24% rest of the world. It is reasonable to suppose that most are not experts in AI & robotics, and that they are unlikely to be publishing in the top venues in AI and robotics like IJCAI, AAAI or ICRA. They are educated (85% have an undergraduate degree or higher), young (more than a third are 34 or under, 59% are under 44 and just 11% are 65 or older), mostly employed or in higher education (more than two thirds are employed and one quarter are in or about to enter higher education) and relatively affluent (40% reported an annual income of $100 000 or more).

The questionnaire itself had 8 questions. The first 7 questions asked respondents to classify 10 occupations from the training set, whilst the last asked for estimates when HLMI might arrive. The first of the eight questions asked for a classification of the 5 occupations most at risk from automation according to Frey and Osborne's classifier as well as the 5 occupations least likely to be at risk. To help respondents, a link was provided next to each occupation describing the work involved and the skills required.

The second of the eight questions in our survey asked for a classification of the next 5 occupations most at risk from automation according to Frey and Osborne's classifier and the next 5 occupations least likely, and so on till the seventh and penultimate question. The final 8th question asked for an estimate of when high level machine intelligence would be reached.

Within each of the first 7 questions, the 10 occupations were presented in a random order. Our intent was to make the initial questions as easy as possible to answer. In this way, we hoped that participants would not give up early, and might be better prepared for the potentially more difficult classifications later in the survey.

The 8th and final question asked for an estimate of when there was a 10%, 50% and 90% chance of HLMI. This repeats the question asked in Müller and Bostrom's survey. The options presented were: 2025, 2030, 2040, 2050, 2075, 2100, after 2100, and never. To compute the median response, we interpolated the cumulative distribution function between the two nearest dates.

5 Results

The results are summarized in Table 1. The experts in robotics were most cautious, predicting a mean and median of 29.0 out of the 70 occupations in the training set at risk from automation (95% confidence interval of 27.0 to 31.0 occupations at risk). The experts in AI were slightly less cautious predicting a mean of 31.1 occupations at risk and a median of 33 (95% confidence interval of 29.6 to 32.6 occupations at risk).

The difference in means between the robotics and AI experts does not appear to be statistically significant. A two-sided student t-test on the number of occupations predicted at risk of automation failed to reject the null hypothesis that the population means were equal at the 95% level (p value, the probability of the observed data given the null hypothesis is true of 0.096).

The non-experts in our survey typically predicted significantly more occupations were at risk of automation than the experts. They predicted a mean of 36.5 occupations at risk of automation and a median of 37 (the 95% confidence interval is from 35.6 to 37.5 occupations at risk).

The differences between the predictions by the non-experts of the number of occupations at risk of automation and those of either the robotics or the AI experts appear

Table 1   Descriptive statistics about number of occupations predicted to be at risk of automation in next two decades. Confidence intervals are at the 95% level.

Group	Sample size (n)	Predicted number of occupations likely at risk of automation (out of 70)			
		Mean	Median	Standard deviation	Confidence interval
Robotics experts	101	29.0	29	10.1	(27.0, 31.0)
AI experts	200	31.1	33	10.8	(29.9, 32.6)
Non-experts	473	36.5	37	10.9	(35.6, 37.5)

to be extremely significant statistically. Two-sided student t-tests rejected the null hypothesis that the population means for the non-experts and the experts in robotics were equal, and the null hypothesis that the population means for the non-experts and the experts in AI were equal (both p values less than 0.0001).

The prediction by the non-experts in our survey of the number of occupations at risk of automation of a median of 37 occupations at risk is identical to the 37 occupations labelled at risk in the original training set in the original Frey and Osborne study.

At the end of the survey, we asked participants to estimate when there was a 10%, 50% and 90% probability of HLMI. This repeats a question asked in the original Müller and Bostrom survey. Also, as in Müller and Bostrom′s survey, we defined HLMI to be when a computer can carry out most human professions at least as well as a typical human.

The results of this question are summarized in Figs. 1 to 4. The robotics and AI experts typically predicted that HLMI was several decades further away than the non-experts. Again, there was little to distinguish between the AI and robotics experts themselves, but they were much more cautious than the non-experts in their predictions. The experts typically predicted HLMI was several decades further away than the non-experts.
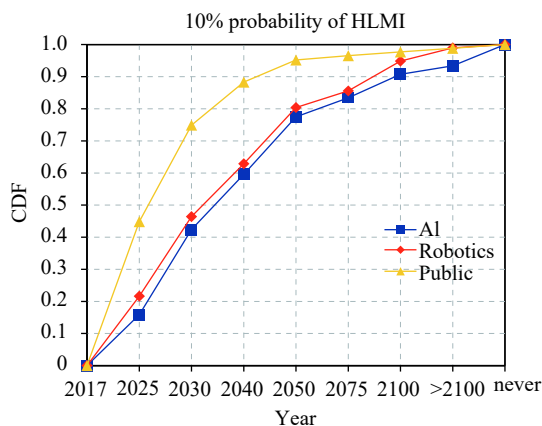

Fig. 1 Cummulative distribution function (CDF) for the prediction of a 10% probability of high level machine intelligence (HLMI). This was defined to be when a computer can carry out most human professions as well as a human.

For a 90% probability of HLMI, the median prediction of the experts in robotics was 2118, and 2109 for the experts in AI. By comparison, the median prediction of the non-experts for a 90% probability of HLMI was just 2060, around half a century earlier.

For a 50% probability of HLMI, the median prediction of the robotics experts was 2065, and 2061 for the AI experts. This compares with the non-experts whose median prediction for a 50% probability of HLMI was 2039, over two decades earlier.
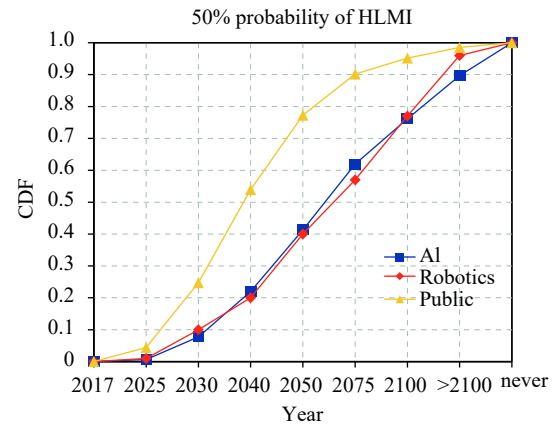

Fig. 2 Cummulative distribution function (CDF) for the prediction of a 50% probability of high level machine intelligence (HLMI).
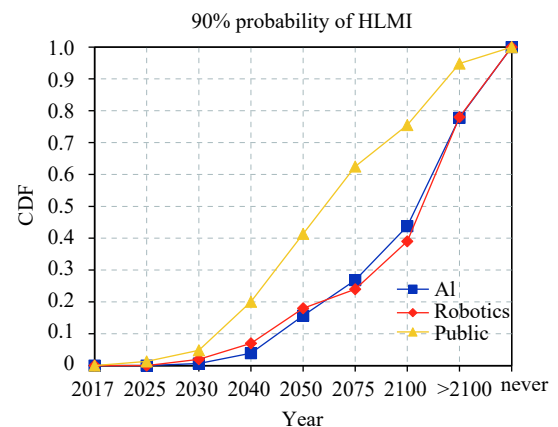

Fig. 3 Cummulative distribution function (CDF) for the prediction of a 90% probability of high level machine intelligence (HLMI).

Finally, for a 10% probability of HLMI, the median prediction of the robotics experts was 2033, and 2035 for the AI experts. By comparison, the median prediction of the non-experts for a 10% probability of HLMI was 2026, nearly a decade earlier.

The predictions for the number of occupations under risk of automation were consistent with the predictions of when HLMI might be achieved. See the clear trend in Fig 4. Respondents who predicted a later date for HLMI typically predicted fewer occupations at risk of automation. Similarly respondents who predicted an earlier date for HLMI typically predicted more occupations at risk of automation.

In summary, the AI and robotics experts typically predicted later dates for HLMI and fewer occupations at risk. On the other hand, the non-experts typically predicted earlier dates for HLMI and more occupations at risk of automation.

The respondents in Müller and Bostrom′s study were closest in their predictions of when HLMI might be achieved to the group of non-experts in our survey. For a
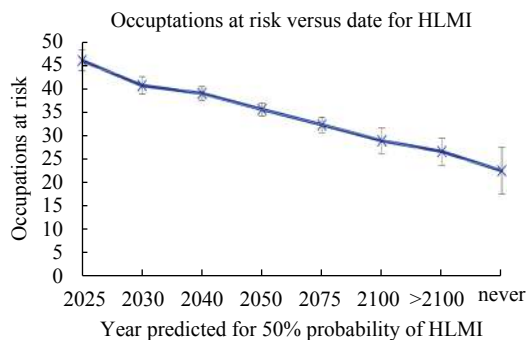
Occupations at risk versus date for HLMI

Fig. 4 Mean number of occupations at risk of automation against year predicted for a 50% probability of high level machine intelligence (HLMI). Error bars give 95% confidence interval.

10% probability of HLMI, Müller and Bostrom's study had a median prediction of 2022, and 2040 for a 50% probability of HLMI. For a 10% probability of HLMI, the non-experts in our study had a median prediction of 2026, and of 2039 for a 50% probability of HLMI. However, for a 90% probability of HLMI, our non-experts were more optimistic than the respondents in Müller and Bostrom's study. The median prediction for a 90% probability for HLMI by the non-experts in our survey was 2060, compared to a median of 2075 in Müller and Bostrom's study.

6 Differences

We looked more closely at the differences between the predictions of the experts and non-experts in our survey, and between the predictions of the experts in our survey and the predictions in Müller and Bostrom's study.

Ironically, given that economists have often been the loudest voices in warning of the risks of technological unemployment, the occupation in our survey on which experts and non-experts most differed was the job of economist. Only 12% of the experts predicted that the job of economist was likely to be automated in the next two decades compared to 39% of the non-experts.

The Bureau of Labour in the U.S. predicts an average 5-9% growth in the number of economists over the next decade. Frey & Osborne's training data classified economist not to be at risk of automation. However, their classifier put the risk of automation at 43%. We would question this prediction. Even if some parts of an economist's job can be automated in the next two decades, we doubt that economists should be too worried about their own technological unemployment.

The next largest difference between experts and non-experts in our survey was for electrical engineer. Only 6% of the experts predicted that the job of electrical engineer was likely to be automated in the next two decades compared to 33% of the non-experts. The Bureau of Labour in the U.S. also predicts 5-9% growth in the number of electrical engineers over the next decade. O*NET

breaks the job down into tasks such as designing electrical instruments, and coordinating manufacturing that are unlikely to be automated soon. Only a few aspects of the job of electrical engineer like technical drawing are likely to be automated in the next two decades. Frey & Osborne's study agrees with this prediction.

The third largest difference between experts and non-experts in our survey was for technical writer. 31% of the experts predicted that the job of technical writer was likely to be automated in the next two decades compared to 54% of the non-experts. The Bureau of Labour in the U.S. actually predicts a faster than average 10%–14% growth in the number of technical writers over the next decade.

Despite computer programs being able to write short news reports, computers still have a long way to go to be able to write long and detailed technical documents. We therefore agree with the experts in our survey in predicting that technical writers should have few fears about technological unemployment. Frey & Osborne's study disagrees such a prediction. Their training data labelled technical writer at risk of automation in the next two decades. And their classifier gave a 89% probability for automation.

The next largest difference between experts and non-experts in our survey was for civil engineer. Only 6% of the experts predicted that the job of civil engineer was likely to be automated in the next two decades compared to 30% of the non-experts. The Bureau of Labour in the U.S. again predicts a faster than average 10%–14% growth in the number of civil engineers over the next decade. As with electrical engineer, we predict that only a few aspects of their job are likely to be automated in the next two decades.

Other occupations where the experts and non-experts in our suvey differed significantly include law clerk, market research analyst, marketing specialist, lawyer, physician and surgeon. In each case, around 20% more non-experts predicted that these jobs were likely to be automated in the next two decades than the experts.

7 Discussion

Our results suggest that experts in robotics and AI are more cautious than non-experts in their prediction of the number of occupations at risk of automation in the next decade or two. The experts in our survey were also more cautious than the training set used in Frey and Osborne's study. This caution can be explained by their expectation that HLMI may take several decades longer than the public expects. We did not find any significant differences between the predictions of the experts in robotics and the experts in AI. Despite being more cautious, both groups of experts still predicted a large fraction of occupations were at risk of automation in the next couple of decades.

There are many other factors that need to be taken into account in deciding the impact that automation might have on employment: We must also take account of the economic growth fueled by productivity gains, the new occupations created by technology, the effects of globalization, changes in demographics and retirement, and much else.

It remains an important open question if there will be an overall net gain or loss of jobs as a result of automation and technological changes. This is clearly a matter that society must seriously consider further. There are many actions possible to reduce the negative impacts of automation. We should, for instance, look to augment rather than replace humans in roles where this is possible.

Even in occupations where humans look set to be displaced, our survey holds out some hope. Whilst the potential disruptions may be large, there could be more time to adapt to them than the public fear. Our study also suggests that more effort needs to be invested in managing the public′s expectation about the rate of progress being made in robotics and AI, and of the many technical obstacles that must be overcome before some occupations can be automated. Robotics and AI remain challenged in several fundamental areas like manipulation, common sense reasoning and natural language understanding. Funding for AI research has suffered "winters" in the past where public expectations did not match actual progress[12]. We should be careful to avoid this in the future.

Acknolwdgements

References

[1] T. Poggio, H. Mhaskar, L. Rosasco, B. Miranda, Q. L. Liao. Why and when can deep-but not shallow-networks avoid the curse of dimensionality: A review. *International Journal of Automation and Computing*, vol. 14, no. 5, pp. 503–519, 2017. DOI: 10.1007/s11633-017-1054-2.

[2] The World Economic Forum. The Future of jobs: Employment, skills and workforce strategy for the fourth industrial revolution. Global Challenge Insight Report, The World Economic Forum, 2016.

[3] J. M. Keynes. Economic possibilities for our grandchildren. *Essays in Persuasion*, J. M. Keynes Ed., New York, USA: Palgrave Macmillan, pp. 321–332, 1932.

[4] C. B. Frey, M. A. Osborne, *The Future of Employment: How Susceptible are Jobs to Computerisation?* Oxford, UK: University Martin School, 2013.

[5] H. Durrant-Whyte, L. McCalman, S. O′Callaghan, A. Reid, D. Steinberg. The impact of computerisation and auto-

[6] C. B. Frey, M. A. Osborne, C. Holmes, E. Rahbari, E. Curmi, R. Garlick, J. Chua, G. Friedlander, P. Chalif, G. McDonald, M. Wilkie. Technology at work v2.0: The future is not what it used to be. Oxford University Martin School, UK, 2016.

[7] V. C. Müller, N. Bostrom. Future progress in artificial intelligence: A survey of expert opinion. *Fundamental Issues of Artificial Intelligence*, V. C. Müller, Ed., Berlin, Germany: Springer, pp. 555–572, 2014.

[8] Y. LeCun, Y. Bengio, G. Hinton. Deep learning. *Nature*, vol. 521, no. 7553, pp. 436–444, 2015. DOI: 10.1038/nature 14539.

[9] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, D. Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, vol. 529, no. 7587, pp. 484–489, 2016. DOI: 10.1038/nature16961.

[10] T. A. Patel, M. Puppala, R. O. Ogunti, J. E. Ensor, T. C. He, J. B. Shewale, D. P. Ankerst, V. G. Kaklamani, A. A. Rodriguez, S. T. C. Wong, J. C. Chang. Correlating mammographic and pathologic findings in clinical decision support using natural language processing and data mining methods. *Cancer*, vol. 123, no. 1, pp. 114–121, 2017. DOI: 10.1002/cncr.v123.1.

[11] O. Etzioni. No, the experts don′t think super intelligent AI is a threat to humanity. *MIT Technology Review*, Massachusetts Institute of Technology, Ed., Cambridge, MA, USA: Massachusetts Institute of Technology, 2016.

[12] J. Hendler. Avoiding another AI winter. *IEEE Intelligent Systems*, vol. 23, no. 2, pp. 2–4, 2008. DOI: 10.1109/MIS. 2008.20.

mation on future employment. *Australia′s Future Workforce?* Chapter 1.4.

Toby Walsh received the B.A. degree in mathematics and theoretical physics from University of Cambridge, England in 1986, and the M. Sc. and Ph. D. degrees from Department of Artificial Intelligence, University of Edinburgh, Scotland in 1987 and 1990 respectively. He was subsequently a researcher in Scotland, France, Italy, England and Ireland. Currently, he is Scientia Professor of artificial intelligence at School of Computer Science and Engineering, University of New South Wales, Australia, as well as group leader at Data61, and guest professor at the Technical University of Berlin, Germany.

His research interests include automated reasoning, constraint programming, propositional satisfisfiability, preference reasoning, social choice, game theory and computational economics. Most recently, he has become interested in the societal impacts of artificial intelligence. He has been elected a Fellow of the Australian Academy of Science, and of the Association for the Advancement of Artificial Intelligence. He has won the Humboldt Research Award and the NSW Premier′s Prize for Excellence in Engineering and ICT.

E-mail: tw@cse.unsw.edu.au (Corresponding author)
ORCID iD: 0000-0003-2998-8668