



Report of the Roundtable on Artificial Intelligence and Employment

**Memorial Room
Jesus College, Oxford**

November 26-27, 2018

Project Authors:

Professor James Maclaurin, Associate Professors Colin Gavaghan and Alistair Knott, Dr John Zerilli
and Joy Liddicoat.

Table of Contents

Background	3
Introduction	3
<i>Uses of AI in the professions</i>	4
<i>The discourse about AI</i>	6
<i>The futures of work and impacts of AI on employment</i>	7
<i>Effects of AI on human workers</i>	8
<i>Gender and AI</i>	9
Regulatory Issues and Challenges	10
<i>Profession specific regulatory responses</i>	10
<i>Ethics</i>	11
<i>Liability and Accountability</i>	11
<i>Human rights</i>	12
<i>Professional Education</i>	13
Emerging regulatory responses	14
<i>Current regulatory approaches</i>	14
<i>Professions' Responses</i>	15
<i>Developing New Regulatory Models</i>	16
Summary	18

Background

What effects will advances in Artificial Intelligence (AI) have on employment? This question looms large in discussions about AI. Public debate focusses on likely effects of AI on human workers: AI systems threaten to eliminate some jobs altogether, and to alter dramatically the nature of other jobs, as parts of existing jobs are automated. However, as AI systems enter the workplace, AI may also have larger-scale effects on whole professions, and the function they perform in society, as well as on individual workers. To examine these effects, in November 2018 we convened a roundtable of experts at Jesus College, Oxford University, to consider case studies of several professions. We were very grateful to Tom Douglas and Hannah Maslen from Jesus College for hosting the roundtable.

Introduction

This roundtable examined several professions that are in the process of being disrupted by AI. Our research focus is, loosely speaking, on ‘Professions’ (with a capital P), such as Law, Journalism, Education and Management, whose practice is instrumental in regulating and organising our society. We are equally interested in the effects of AI on human workers in these professions, and its effects on the professions themselves. At the roundtable, we focussed on what changes to policy and regulation (if any) might be needed to respond to both types of effects. We sought to identify trends and issues in AI and employment and to understand emerging responses, including what, if any, policy or regulatory responses may be needed. AI technologies pose fascinating legal, practical and ethical challenges, which require interdisciplinary solutions. Our international experts came from a variety of jurisdictions including the United Kingdom, Europe, the United States, Canada, South Africa and Brazil. Participants were from academic institutions, industry, civil society, the technical community, and professional bodies. Participants shared insights into, and critically examined, the current and likely future impact of AI on the law, education, journalism and corporate management professions.¹

This report summarises the roundtable discussion including: the benefits of AI, the nature of the discourse about AI, the futures of work and specific impacts of AI on employment, how workers view the effects of AI on their professions and employment, and the gendered impacts of AI on employment and definitions of work. A summary of regulatory implications follows. Finally, there is a brief critique of current regulatory approaches and how professions are responding along with some emerging regulatory models.

The roundtable was held under Chatham House Rules, where contributions were anonymous but not confidential. In this report we have referred in general to “participants” or “views”, but emphasise that these terms do not reflect any particular individual’s views nor do these necessarily indicate consensus. Any errors or omissions are ours alone.

¹ The full workshop programme and participants details are available at: <https://www.cs.otago.ac.nz/research/ai/AI-Law/Workshop2018/index.html>

AI and Employment in the Professions

AI and machine learning technologies are affecting each of professions in different ways: from creating new legal services to offer clients and creating new efficiencies in legal practice; to creating new forms of journalism and news media story generation; to enabling new forms of learner-centred personalised education and enhancements to teacher development; and to new forms of management and corporate accountability. All participants anticipated these effects to grow, raise new issues and to have significant impacts on their professions, but had differing views on the entry points, pace and timing of change. Some cross-cutting issues were identified and are briefly outlined below.

Uses of AI in the professions

Natural language processing and other AI tools are being used to develop diverse new services. For example, these tools are being used to enable a much wider range of people to have access to legal information at a fraction of current costs. Inadequate access to legal information is a significant barrier to access to justice and these new tools have significant potential to improve access to justice by directly responding to legal questions without the need for a lawyer. New forms of contract are being developed, such as inter-active contracts, which would enable consumers to better understand their insurance policies. Contract analytics was a rapidly developing field with a diverse range of legal products available. For example, products that can create a digital contract summary (through digital review of lengthy contractual documents), a preliminary assessment of governance and risk issues, recommendations for triage, work-flow and prioritisation as well as draft reports and suggested benchmarking for progress. Other products enable speedy analysis of millions of pages of documents to speed up investigation and pre-trial evidential processes.

While the use of AI in legal practice is relatively new, “AI for Education” has been developing for more than 40 years. Now more commonly referred to as “AI in Learning Environments”, the area has more history than in law and journalism. Intelligent tutoring systems, for example, are moving learning beyond a simple dichotomy of ‘teacher *or* robot’ to ‘teacher *and* robot’. New tools, such as Tardis and Echoes, can allow better understanding of how teachers and students learn, enhance teacher practice and change the nature of learning, but only, it was emphasised, if the unique benefits and respective strengths of artificial and human intelligence were well understood and applied.

In journalism, AI was being used across the realm of news production including editorial functions such as news gathering, content production, news verification, and moderation (for example, to moderate comments on news stories online) and for business functions (including distribution and subscription services). Task automation technologies were being used, for example, to classify data, identify trends, and ‘find needles in haystacks’ of information, to automate routine editorial functions and routine reporting, together with production of summaries and generate video from text summaries. Similar technologies were being used to develop interactive customer services such as

chat bots, automated newsletters, flexible payment options and to segment large customer bases into diverse parts to facilitate more differentiated, pay-for-use, services. These technologies were similar to those used to develop more heavily criticised forms of media content such as 'fake news' and unsolicited profiling by online social media platforms. In addition, new journalistic skills have developed to enable, for example, creation of local news content through programming and use of open data to pull information and automatically generate individualised content about dozens of cities or towns, which can be offered to news media outlets in those cities and towns.

The value proposition for each use of AI varied. In journalism, this ranged from augmenting reporter capabilities (for example, to deeper analysis more quickly), to creating shifts in tasks to enable editors to do higher value work or developing better local or personalised content for users to increase the news media outlet's exposure and options for monetizing content.

In relation to AI and corporate liability and applications to corporate management, participants heard of the rise of the 'algorithmic boss', with a case study examining the gig economy, such as Amazon Mechanical Turk a platform which enabled requesters to contract workers for micro jobs (such as image detection and object recognition). On the one hand this new way of working enabled those undertaking the work to participate in micro-entrepreneurship, providing flexibility to choose when and where to work. On the other hand, this was hiding work, creating a micro task platform which enabled requesters to tightly control work insecurity. For example, in the gig economy a worker does not get sacked, they get 'de-activated'. This represented a movement back to a model of work that is focused on discrete tasks rather than jobs.

In this context, micro task platforms such as Amazon Mechanical Turk was described as 'artificial artificial intelligence', because these were effectively selling humans as a service (Bezos) and rendering the work and workers invisible (see critique of Deliveroo ad). This was similar to the original mechanical Turk, which had a human chess player concealed in a box. In addition, wages appear to be extremely low and with gender bias (Abigail Adams). There was a strong caution against technology determinism in relation to the future of work, there being, instead, no such thing as 'the future of work', but in fact many possible futures of work.

Participants were introduced to three problems with regard to AI application to management tasks. The first was a significant miss-selling or over-promising of AI and what it is, there being no 'general AI'. Rather most was 'narrow AI', simply, statistical models which have issues relating to design, data, and interpretation of results. This raised normative legal issues, and some bright lines were needed in the definitions of AI. Secondly, the notion of 'Gen 2.0 automation' and very disparate estimates of size and rates of human job displacement and diverse proposals for how corporates should respond. Thirdly, there were consent and transparency issues, particularly with human oversight and decision atrophy.

In relation to AI and corporate governance, some new issues were outlined. For example, the rise of the so-called AI corporation such as the Adidas Speedfactory, new oilrigs (which have only 5 people working on them) and Dulux (which has just 150 maintenance workers compared to the previous 800). However, it was also noted that some of those worried about the implications are those in middle management positions, as fewer workers mean fewer managers are needed. Concerns were outlined about the models of technology corporate governance, with reference to the 13 major technology companies which are all owner controlled and do not meet good governance standards for publicly listed companies. This created significant risk for investors and for the public. Another issue was whether and if so how to develop legal personhood. For example, in the UK human rights “aspirational” but what is the best gateway to rights (human rights V State, Realist V Contractarian rights. Finally, the emergence of genuine enhanced decisions, for example, humans making decisions with assistance of driverless cars and the moral choices and dilemmas this was creating.

The discourse about AI

The nature of the discourse about AI was a factor in how AI was affecting the professions and discussion about AI and employment. The dystopian or utopian futures binary was prevalent, as was framing development and deployment of AI as robots *against* humans or robots displacing humans. For example, a dystopian narrative of humans being displaced by machines created a climate of fear which made it difficult for professions to openly discover the risks and opportunities of new technology. As a result, they risk feeling the effects of the new technology without adequate preparation. Such an approach also inhibits developers and regulators taking a more nuanced approach to discussion about uses and effects, particularly where these are quite different in different professions.

Participants cautioned against presenting a technology as “AI” when it is not AI, highlighting, for example, the ‘miss-selling’ of AI which, it was suggested, was simply using statistical models for design, data and interpretation. The question of definition (“what is AI”) was important because it also raised a normative legal problem and related issues such as how to explain or unpack decisions.

Similar concerns about the AI discourse arise in relation to description of the so-called ‘benefits’ of AI. For example, saying that AI tools enable companies or government agencies to “get rid of waste”: this kind of language could appear to be professional disparagement of working class jobs. Conversely, it was noted that some professional work, particularly early in a career could be described as not meaningful (such as routine labelling of documents for evidential purposes in pre-trial proceedings) and that not being able to describe the work as lacking meaning might have an impact on professional career decisions.

New terminology might be needed for new forms of work which are emerging, forms of work that are not yet readily visible both legally and to users of the services that have been developed with AI (such as image identification and processing). Due care and more critique about use of the term “AI” was encouraged, including consideration of who is using AI, where and particularly to consider government use where public money is used to procure services from those companies using or

developing AI. A further reason identified for the need for care with terminology was users or clients may have unrealistic expectations that their professional advisers are using AI, when they are not. Disaggregation of the technology and its use was suggested, for example, to let clients know when AI is being used and to help shape and better use the tools themselves. The gendered nature of the AI discourse was also noted (see further below).

The futures of work and impacts of AI on employment

Participants shared research and insights into macro-economic research on the historic impact of computing on jobs and weighed evidence of likely future impact of AI on employment. Declining costs of computing, which was increasingly doing more routine tasks was considered, along with information about increasing job polarisation between low, middle and high wage earners in developing countries.

Participants noted there have been many different reports and studies of the ‘future of work’ in relation to AI, with many predictions, yet there are many possible futures depending upon how the technology is developed and deployed. There were dangers therefore in narrowing too quickly likely future scenarios. Low skilled or highly routine jobs were identified as relatively exposed to automation, but displacement effects also give rise to new job opportunities. One conclusion was the new technologies would have both positive and negative effects on work, including job creation and job loss. Areas where humans would retain comparative advantage over AI included jobs that require social intelligence, creativity, perception, and manipulation. New forms of marginalisation and poverty along with new forms of work will emerge.

Participants discussed the implications of AI on the nature of employment regulation, including recruitment, redundancy, human oversight of AI, transparency and liability for AI decision-making. The impact of automation and artificial intelligence on the employment relationship in light of basic tenets of industrial relations (technology of production, the employment relationship and issues that are important in labour law) was highlighted. The impact of technologies on the organisation of production which displaces employees was also considered (for example, the rise of subcontracting and outsourcing and the emergence of the “sharing economy”, “worker on demand” apps and micro jobs, such as Amazon Mechanical Turk, Crowdfunder and Click Worker).

These emerging AI related economies gave rise to legal and regulatory issues. For example, in classification of employment contexts, new and complex definitions of employer and employee, problematic regulatory arbitrage, concerns about management by algorithm, job polarisation and the extra-territorial application of labour laws where workers reside in countries outside the jurisdiction of the employer. Participants also reflected on these issues in the broader context of ILO research on the future of work including new business models and economies, changes in the nature of work and jobs of the future.

Participants noted many of the changes were happening within existing regulatory frameworks and that while these frameworks were not yet considered wholly problematic, some gaps were emerging. For example, there may be gaps in laws on redundancy which require new forms of redundancy in labour law where humans are working alongside AI but have only parts of job tasks, rather than whole jobs, made redundant. Further work on regulatory responses to redundancy and replacement issues, including new types of responses that take into account the changing nature of work, would be needed.

Participants broadly agreed that while the impacts of AI may include technological unemployment (the prospect that people will lose their jobs to intelligent machines), the more likely scenario is that humans are, and will increasingly be, working with AI in a variety of ways. Participants considered changes to human jobs will, in general, happen iteratively, rather than in single acts of wholesale disruption but tipping points might be reached in some cases which appear in a more disruptive manner. This theme, of machines and humans working together was not a new one, nor peculiar to AI. Human work that was augmented by AI (“AI augmented human work”) was considered the more likely next step in light of how natural language processing and other forms of AI and their related uses are already developing.

Effects of AI on human workers

The roundtable focussed on the effects of AI on human workers in the professions as well as the effects on the professions themselves. In journalism, the answer to whether journalists are satisfied with new ways of working was not clear. Some new forms of work meant different forms of satisfaction (for example, writing 200 stories in a week or seeing their story on another agency’s news site). New challenges had emerged such as a loss of journalistic style or the need to re-write initial products if the generated stories were inadequate. There were related questions about whether the use of AI to create a story should be made clear or whether the story should be published under the journalist’s name. It appeared that some publishers had concerns about identifying auto-generated content, preferring to name the journalist as author (because the reporter decides what story will be written, generates the data questions and writes the story template).

A particular issue for the professions is that some people they serve may not ‘trust’ AI, but may rely on the regulated nature of a profession to give them confidence that AI technology is being used in a trustworthy manner. A recent survey about attitudes to working with AI was referred to, which found young people more optimistic than older people about AI, but more concerned about the impact of AI on inequality than some people in older age groups.

Caution and care was recommended in the way in which AI and human work is described, particularly professional human work. Some considered too much discussion about the impact of AI on work is between intellectual workers and not with those who are most at risk. It was important not to patronise those workers, to listen to how people feel about work and the impact of AI on their

employment and to bear in mind that most people who do routine work enjoy their jobs, value work and are devastated if they are unemployed.

Gender and AI

The issue of AI and employment was considered in light of the already significant gendered impacts of digital technology. For example, participants were referred to the Malala Fund research on women and girls use of technology, which found that women and girls are not using technology to find work but for more social connection purposes. There was a need to consider the implications of this for AI and employment research and in the light of the overarching inequalities in women's digital skills and access to technologies. Against this background, there were concerns that new technologies may increase gender inequalities in both formal and informal work. For example, while new forms of employment may emerge, women may be disproportionately affected by displacement given their higher representation in jobs involving routine tasks. The connection between technological skills and attachment to work was also highlighted as a risk for increasing gender inequalities, particularly for older women who may not be able to access tertiary education and may rely on updating technology skills through learning in the workplace.

The importance of examining who is building the AI was emphasised as a means to avoid gender bias. For example, if the goal of an AI tool is to achieve good development outcomes for women, women need to help build the AI tool. Concerns were also raised about the potential negative impact of AI on women's unpaid work and the strong inverse link between unpaid work and economic empowerment. In relation to AI and employment in the professions, it was noted there are significant gender inequalities in the professions as well as differences between groups of women workers (such as professional women, compared to care workers compared to those in manual jobs).

One suggested response was to try to change the narrative about AI and rather than developing policy by template, to consider more deeply research on digital literacy and better understand disaggregated data on different groups needs. The voices of the women, of others who are the most vulnerable to being displaced and of those working in the informal sector were missing from the current discourse and needed to be included. A suggested alternative was to frame the AI discourse in relation to the right to development. The United Nations Sustainable Development Goals (SDGs) were offered as good policy reference point for this as all governments have signed up these at the United Nations. Other options were to consider new ways to engage people in the discourse about AI. For example the "Just transition" / "the Green New Deal" initiatives, were using exciting ways to engage people with humour (for example, around "fully automated luxury communism" or "unemployment for all").

Participants also considered the future of unwaged work which, it was suggested, should be added to consideration of what is "work". Questions in relation to household work included: how it is changing and which domestic tasks are changing; who is designing domestic work apps as opposed

to who is consuming these? Many housework apps simply aim to make housework “fun” or “more enjoyable” (for example by having music choices for tasks) rather than changing the nature of that work. The gendered nature of technical skills in the home was also a factor in how women’s domestic work might be affected by automation of household tasks. A further question was what new tasks might emerge to fill the purported time gap? For example, automatic washing machines changed expectations of cleanliness so that people spend the same amount of time on washing, but in smaller, regular amounts. This area raised interesting questions of how the future of work might impact on the future of home life in light of its gendered nature.

Regulatory Issues and Challenges

Participants noted that regulatory responses and related policy work will need to take into account that small, iterative changes to how AI will affect work may lead to tipping points of wider employment disruption. For example, the public discourse on AI and machine learning was largely negative and dystopian and this was negatively affecting the ability have more critical and nuanced discussion about AI, which was considered essential for good policy and regulation. This discourse might not only stifle good policy and regulatory development, it might also result in calls for significant regulatory responses out of step with how the technology was developing. Another element to the discourse was to counter the notion that regulation was a barrier to innovation: regulation can be a driver of innovation and there was nothing remotely luddite about making technology use responsible.

Profession specific regulatory responses

Discussion focused on both the effects of AI on human workers in professions and its effects on the professions themselves, including what changes to policy and regulation (if any) might be needed to respond to both types of effects. A strong theme was that regulatory responses should be sector specific, rather than general, for several reasons.

First, the new technologies are being used for very different purposes in different professional contexts (for example, law compared to education, compared to journalism compared to management). Second, even within each sector, the technologies were being adopted at very different speeds, in diverse ways and at different levels, with no uniform adoption or approach. Third, there is no ‘one sized fits all’ approach to regulating these professions now and it was unlikely a general regulatory approach to professional regulation specifically for AI would be workable. In the area of education, for example, the existing education system and related policies were identified as a challenge and changing these would take considerable time, whereas in journalism or law change might happen more quickly.

In relation to regulation of work more generally, participants noted ILO research in 121 countries found that most people want to work. This suggested that the regulatory focus and policy debate should not be about whether or not people work but “are we getting better work”. New forms of

work may require new forms of regulation or amendment of existing intellectual property models. For example, workers in the global south doing image assessments for machine learning tools raises questions of whether new forms of intellectual property law be needed to determine who owns the rights in a new service that has been built with the assistance of thousands of such workers spread in many diverse settings.

Ethics

The ethics of AI systems need urgent and systemic discussion, regarding the more obvious examples about their possible use in war or for cyberattacks. Less obvious, but equally compelling areas for ethical discussion were emerging, for example, in relation to use of AI for legal information services, a significant risk is that such tools will stifle the common law. The different regulatory frameworks for law practice in the United Kingdom compared to the United States of America were noted. In the United Kingdom a lawyer owed ethical obligations first to the court (as an officer of the court) and second to clients. Ethical guidance was therefore needed on matters such as non-disclosure agreements or whether to release code to clients (and, if so, in what form). These issues also related to concerns about transparency and understandability and were in tension with concerns about protecting intellectual property.

Another issue was with ensuring the technology stays current, that issues of bias in underlying data are addressed through regular data updating so any errors do not affect the most vulnerable. Some considered the lack of universal standards for data science makes it difficult for those outside data science disciplines to have confidence in the use of data sets.

Other issues noted were whether the ethical standards for those in different professions should be the same or different, for what purposes and whether these should also be different for government and business.

Liability and Accountability

If AI is just another technological tool, the first regulatory step is some stress testing of the current laws on liability and accountability first, to determine whether a new approach to liability and/or accountability needed. For example, if the purpose of liability and accountability laws is to prevent harm or bias, there were some existing provisions which may apply such as consumer protection law, privacy or human rights laws, violations of professional codes of conduct or others, depending on the particular circumstances and the nature of harm that has occurred and its cause.

Participants considered assigning liability will be difficult in some situations under current law, including how to deal with liability for bias. In journalism there were tensions, for example, between intellectual property and fair use. The recent global attention on issues of accountability for content and distribution, along with related difficult questions of liability for 'fake news' on new platforms that do not hold themselves to the same standard of liability were creating professional difficulties.

For example, new forms of content distribution pushed liability to users, contrary to journalistic professional standards. When writing news stories with new technologies, other journalistic standards may need fresh consideration, for example, where the journalistic source is an open data set.

Examples of litigation of algorithms under current laws were discussed. In these, it was important to understand the context in which AI is deployed, for example, in the case of Houston teachers. One option, considered by the European Committee, was the creation of legal or electronic personhood, namely, to create a legal fiction that might be applicable, similar to concepts of corporate liability. Participants noted that other new concepts of liability and personhood might emerge, drawing an analogy with, for example, New Zealand law which has recently conferred legal personhood on a river. Conferring personhood on an algorithm was considered challenging.

Tort law and consumer protection laws were considered, along with warranties, which appeared to be a grey area. For example, if a customer is allowed to assist in training the machine learning tool, or a client is helping develop it, what is the legal situation if a client wants to offer the new tool back to their community - what would be under warranty and what might the developers responsibilities be? Leaving the matter to insurance law was not considered adequate because, for example, there may be non-financial consequences or contracts might include unenforceable provisions (such as anti-competitive clauses).

A further complexity was the need to align regulatory provisions with the nature of ethnical and other tools to maintain the technology, for example, introducing legal liability for errors. In the development area a culture of error reporting was to be encouraged, for example, to produce validation flags that could correct outcomes or help improve the tool. Introducing legal liability for errors could stifle the ability to uncover mistakes that could be remedied, increasing, rather than decreasing, the risk of harm.

Human rights

The topic of AI and employment raised a number of human rights issues including the right to privacy. Participants noted that while data protection laws were important in the context of AI and data use, these new technologies were raising a much wider range of human rights issues for example, employment rights, freedom from discrimination, freedom of expression and freedom of association.

In relation to employment, AI and online content, concerns were expressed about “deep fakes” and what, in an employment setting, might be the evidential standards for verification of systems that were being used. Similarly, increased and more specific targeting was making accountability difficult (for example, in relation to political advertisements). On the right to work, participants discussed which regulatory issue needs addressing. For example, there is currently a legal distinction between displacement by another cheaper person compared to being made redundant. In the context of the

right to work, where an AI tool can do the same job more cheaply, there are unresolved legal questions as to whether the human role is legally redundant. Nor do regulators appear to have all the answers (in the UK, for example, the ICO has hired a technical team to do algorithmic auditing).

More generally, the global nature of business models suggested that cross-jurisdictional employment practices would need to develop. A brief review of changes in business models indicated some companies were down-sizing their out-sourced businesses in more globally remote countries and instead near-shoring, for example, UK technology companies near-shoring in Belfast. Questions have arisen about the appropriate employment law to apply as well as the application of professional standards across jurisdictions.

In relation to freedom of expression, the recent report of the Special Rapporteur on Freedom of Opinion and Expression was noted, including his assessment that curation of news content by AI may violate the freedom to receive information. A further concern was the vulnerability of voices of dissent and criticism, for example in journalism and news story generation and who makes decisions about whether a particular decision is editorial oversight of content and standards of content or amounts to censorship of dissent or criticism. Some considered that it was harder for corporations to take a more nuanced approach to freedom of expression and content regulation in this context as the moderation tools are not well developed. The New York Times, for example, uses AI tools to assist with comment moderation, but also has editors who moderate comments' rules.

Finally, it was noted that most professions members have professional associations or trade unions, and that unions are also regulated. However participants noted the changes in unions and union culture in the last two decades has resulted in the emergence of significant numbers of workers who do not have experience of joining a trade union. Understanding what steps, if any, trade unions are taking in relation to AI and employment would be helpful.

Professional Education

Participants reflected on the implications of some of these developments on professional education and the readiness of business, educators and profession regulators. In Belfast, for example, near-shoring of corporations has resulted in some law firms doing short term training of new employees, such as employing students with legal training and some technical skills, but law firms are finding legal education is inadequate. Questions are being asked about the roles and responsibilities of universities to better prepare students, rather than relying on law firms to upskill. At the same time, legal educators were not sure how to prepare for some of the new developments. Some considered communication between the two is key and that universities have a role to play and law firms have a role in supporting collateral learning. Bearing in mind the machine learning tools are now the worst they will be, some participants considered a more radical approach to change professional education was needed, but some had experienced strong resistance to this change from academics.

In education, it was noted that AI can be used to inspect the current systems and to expose where the flaws are and this could be very helpful to individual professional teachers as well as to regulators seeking to improve their practice. However, this was not happening at the moment and the same appeared to be true for lawyers and journalists. There were also professional education challenges for teachers who struggled to keep pace with technological change and learn the skills to use AI in their work.

On the question of the specific skills that might be needed emphasis was put on skills that would add value to tools and support augmentation of work (such as philosophy or ethics), rather than replacement or attempting to compete with AI (such as learning to code). On the other hand, while coding was not considered critical, understanding code and how it worked was considered an important skill. In journalism, some new technical skills are needed, but the ability to write stories remains central to journalistic practice.

Emerging regulatory responses

Current regulatory approaches

Participants considered a range of professions and noted their diverse regulatory contexts. Some considered the future of work was especially important for the professions, where regulatory reform can take many years. Others considered there is “no such thing at the future of work” but rather that, as humans have agency over technology, for example, to make work safer rather than using technology to monitor workers, there are many multiple futures of work.

Concerns were expressed about the current libertarian regulatory approach to technology characterised as “permissionless innovation” (which proposes no or minimal regulatory barriers to developing technology), which some considered a regulatory impunity model. It was noted, for example, that almost all other sectors have a presumption of regulatory precaution but the technology sector has refused to accept this: was this also the case with regulation of AI and machine learning technologies?

In response to concerns about permissionless innovation, there was discussion and critique of the role of regulation in innovation. The view that regulation stifles innovation was argued against in favour of the alternative view that regulation is a driver of innovation (for example, some considered the GDPR has potential to drive innovation). Some case studies of consequences from lack of regulation were also noted including the global financial crisis, child labour, modern slavery, lead in petrol, and asbestos use.

In contrast to the permissionless innovation approach, participants noted there appeared to be more willingness from corporations to consider a regulatory approach now because corporations wanted to be able to protect their current and future investments in new AI tools. There are challenges,

however, for example, does thinking through regulatory approaches to autonomous vehicles now make it far sighted or will it be outdated before it comes into force? Participants encouraged a focus on what harms could arise, slowing down and doing deeper reviews to better understand what to keep, change and what to avoid.

Concerns were expressed that the GDPR was becoming a global standard by default and that five US and three China firms also setting AI related standards by default. However, according to the British Standards Institute for good governance none of these companies meet the standards for good corporate governance. This was an issue also for the structure of different professional bodies and the forms of legal or medical practice that might be permitted. A final question was whether regulatory approaches should differentiate between development and deployment.

The point was also made that some countries' do not want to regulate AI because of "security concerns" or because they do not wish to suffer first mover disadvantage by limiting what can be done with the technology. While it is important not to conflate the AI and Cyberwarfare discourse, participants considered it was important to take account of this pressure not to regulate as the AI discourse is taking place in that wider discourse of permissionless innovation.

The interface of administrative rules and private decision-making in this new area was discussed, namely, whether to continue to apply the existing systems such as creating legal safe harbours to hold private party responsible if AI causes harm. Others considered safe harbours a red flag that something untoward was happening in the liability chain. Instead, a broad functional approach might be a new way to exercise control (i.e. if a tool delivers a racist outcome then it is irrelevant how it did so).

Some considered the related definitional issue (what is AI) could be avoided by focusing not on whether or not a particular technology was AI or not, but instead looking at how the technology was used. For example, regulatory lessons from telecommunications changes in the past could be a useful resource.

Professions' Responses

Some professions, and some professional member organisations, have a significant focus on technology and others do not. The England and Wales Law Society, for example, has established a technology and Law Commission with a significant focus on technology and 21st century practice that will report in mid-2019. The Society appointed three Commissioners and is holding hearings looking at AI and its implications for law, human rights and freedoms, and specific issues such as liability. Emerging findings include:

- The safeguards of the GDPR should, but do not, apply to all algorithmic tools (for example, those used in policing escape the safeguards of GDPR);

- Prediction tools are likely to have the most immediate impact on the justice system and will have short term as well as longer term impacts. For example, in the longer term there is a risk the common law will not receive outlying cases and, if it does not, will not develop as it has in the past. This could result in the common law “seizing up” in some areas because new, innovative cases are not brought to the courts. The result is that widespread use of algorithmic tools could stifle the common law;
- AI tools are being developed for commercial purposes rather than public service or public interest purposes and government should consider having more of a role in both development and deployment as well as standard setting and guidance.

Developing New Regulatory Models

Participants observed that just as a multi-disciplinary approach to AI research was critical, so too a multi-disciplinary approach is needed to regulatory models for AI. A single regulator approach was not adequate or appropriate. Gaps in the regulatory frameworks were emerging and new ethical principles, or restatements of existing ethical standards, were one way to guide practice until new regulatory responses emerged, or to test whether additional regulatory responses were needed. Regulators were lining up on different parts of the regulatory spectrum with the World Bank, for example, suggesting lower employment protections may be needed to respond to automation of work and other researchers suggesting the need for stronger protections to deal with these effects.

While surveys show the public are quite positive about aspects of new technologies and AI, they also have worries about inequalities and how benefits can be spread fairly. Supporting policy makers to understand the technology issues was important and initiatives such as Technology Education Project aimed to break down barriers to policy making. However, the rate of change of politicians is far slower than technological change and it was important to consider how best to education policy makers who may not have the skills to engage in some aspects of the debates. This is especially important for the professions which rely on policy makers for regulation or updating regulatory frameworks.

In addition, the pace of change is a factor, in that by the time there is a policy response the technology has changed. New and innovative ways of regulating may be needed such as adaptive policy and instruments which adapt as contexts change: regulation may need to be done differently. One suggested response was to adopt a regulatory model similar to that used for licensing pharmaceuticals, and to have a regulatory body that was able to consider harm and, for example, recommend withdrawal from the market in some circumstances.

In education, for example, it was important to see through the digital veil and consider the social impact of using AI to understand how humans are working with the technology and politicians were not the only ones that could, or should, determine this. A further concern was a regulatory approach of technological solutionism or determinism. Some noted there was an evangelical, rather than purposeful, tenor to some of the AI policy discussion with different governments vying to be the

safest, easiest or cheapest place to develop AI, or that AI will make a country great, greater or great again. These more industrial policies, it was suggested, are flawed for the ICT world and will also be for the AI and employment area for several reasons:

- (a) The view that there is “nice IT - ie that IT is benign and can be trusted to deliver benefit” was fallacious. Rather, these are enabling technologies which enable everyone, good and bad, and there will be unanticipated outcomes, some of which will be significant (eg filter bubbles as an outcome of social media). Policy must therefore be built around harm rather than benign or malicious use.
- (b) Policy is not currently fit for purpose because it is not holistic (too often technology policy is made by techies and in ways that do not suit others).
- (c) There is too much faith in a model of permissionless innovation. Regulation through democratic processes shapes better policy.

A specific challenge for regulation in relation to the effects of AI on employment, including the professions, will be during the transition to AI augmented human work. While new jobs will emerge over time, there are significant issues in the transition period. Job change estimates are influenced by a range of factors but disaggregation, including by gender, is important and policy makers could be at risk if they make decisions based on binary projections. Governments may therefore need to prepare for poorer outcomes than expected during the transition period. This might also suggest a need for longer term policy thinking.

The broad range of possible solutions that are being proposed was outlined. The first range included: increasing public understanding of what’s happening so there are more democratic forms of consent to new technologies; imposing new obligations on corporates under existing law, for example, to re-train workers affected by automation; create new forms of ownership of AI which recognise the mutual ownership with human workers (such as those training AI); introduce a robot AI registration tax (to ease the cost of the transition to re-skilling); introduce universal basic income; and develop new ethical principles (some companies and multi-stakeholder collaborators were doing this already).

A second range was more specific to the governance and regulation of the technology industry. Examples included addressing the corporate governance models issues with owner/founder controlled companies (for example, if these companies do not act in the public interest do not let them set AI policy and setting some governing principles (for example, that AI be used for tasks that are dangerous to humans or that humans can’t do or find difficult). Options to address the ownership concentration of large tech firms included creating public interest AI representatives at board level on large Tech firms and setting responsibilities for AI standards and risk. Another option was to insert normative legal values in statistical models (such as equality, race, gender etc) so that AI is designed to produce the outcome society wants and the law requires, rather than what the AI has made a raw

data decision on. Training users and educating consumers about AI statistical modelling and the problems of outliers, was also suggested, along with better forms of disclosure.

A specific proposal was to treat AI like a pharmaceutical product and create an AI Algorithm Regulator similar to NICE and FDA to enforce public interest standards in automated statistical model decision making. In this type of model the regulator would license AI for use only where it meets quality and human enhancement standards, rather than a human replacement test.

Another view was that given that digital dividends of other forms of ICTs have delivered inequality, we should expect the same will happen with the advent of AI: loss of job security, weakening job rights, loss of fringe benefits (eg leave, pension) and so on. It is not clear that the benefits of job displacement will be spread fairly: some workers may lose rights that took decades to develop and were entrenched in law and which have meaning in society and this needs to be included in AI policy and regulatory thinking.

There was scepticism about growth in demand for caring and creative work as a result of increased use of AI for other work. One suggestion was that if these are expected to be areas where human jobs will increase, it would be expected that government strategies would be focused on non-digital rather than digital skills training, yet no national strategy appears to focus in that area.

Summary

This report summarises issues which emerged from the roundtable discussion including: benefits of AI, the emergence of AI augmented human work, the impact on uptake of AI from the nature of the discourse about AI, specific impacts of AI on employment (including how workers view the effects of AI on their professions and professional education), liability and accountability challenges, cross-jurisdictional and related matters. There are a number of regulatory implications, possible new models and areas for particular attention whichever regulatory or policy approach is taken (including retention of the status quo).

We were grateful to all participants some of whom reflected that the roundtable discussion was not typical of other workshops on AI in that there was a plurality of voices and discussion of a more cautionary approach. We will consider this feedback as well as the implications of the discussion for our future research on AI and employment and the professions.