



Report of the Roundtable

on

Artificial Intelligence And Employment

September 26-27, 2019

University of Otago

Dunedin

Project Authors

Professor James Maclaurin, Associate Professors Colin Gavaghan and
Alistair Knott and Joy Liddicoat

Table of Contents

Background	4
Introduction	4
<i>AI in hiring and recruitment.....</i>	<i>6</i>
<i>AI and employment law</i>	<i>8</i>
<i>Employment and the gig economy.....</i>	<i>12</i>
<i>Ethical and regulatory challenges</i>	<i>13</i>
<i>AI in the legal profession.....</i>	<i>15</i>
<i>AI and health care</i>	<i>16</i>
<i>Digital assistants and counsellors.....</i>	<i>19</i>
Conclusion.....	21
Appendix A: Roundtable Programme and Participants	24

Background

What effects will advances in Artificial Intelligence (AI) have on employment in New Zealand? This question looms large in local media discussions about AI, and increasingly preoccupies New Zealanders preparing to enter the workforce. Public debate focuses on the likely effects of AI on human workers. AI systems threaten to eliminate some jobs altogether, and to dramatically alter the nature of others, as parts of existing jobs are automated.

However, as AI systems enter the workplace, they 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, this Dunedin workshop considered employment related issues alongside case studies of medicine and law and two professions with a more casual, mobile workforce - gig economy workers and personal assistants. We are equally interested in the effects of AI on human workers in these professions, and its effects on the professions themselves. Our focus in this roundtable was on what changes to New Zealand employment policy and regulation (if any) might be needed to respond to both types of effects. This was the fourth and final roundtable in the New Zealand Law Foundation funded project and provides a critical contribution to project direction in its final year.

Introduction

Public debate about AI and employment focuses on the likely effects of AI on human workers. AI systems threaten to eliminate some jobs altogether, and to dramatically alter the nature of others, as parts of existing jobs are automated. However, as AI systems enter the workplace, they 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, this workshop built on the one we convened in Uehiro College, Oxford University in 2018 to consider employment related issues alongside case studies of several professions. In the Dunedin workshop we brought together practising employment lawyers, academic AI and law researchers, policy experts and trade unionists together with people working in the New Zealand AI industry, and from representatives of professions that are being disrupted by AI. Our participants came from a variety of sectors including academic disciplines, private sector, 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 medicine, law and human resources management professions.¹

This report summarises the roundtable discussion including of the uses of AI in hiring and recruitment, AI and employment law and the gig economy and related ethical and regulatory issues. The roundtable was held under Chatham House Rules, whereby 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 attached as appendices are also available at: <https://www.cs.otago.ac.nz/research/ai/AI-Law/Workshop2018/Dunedin-programme.htm>

AI in hiring and recruitment

Participants were introduced to the topic of AI in hiring and recruitment through presentations by James Maclaurin, Colin Gavaghan and Alistair Knott. James Maclaurin provided context for the wider discussion of AI and questions as to its possible effects on employment. Alistair Knott outlined that AI systems can be used in hiring and recruitment to process candidates' textual submissions (CVs and cover letters), conduct interviews using chatbots, scrutinise candidates' nonverbal behaviours in interviews and to process text and video from public social media pages can also be processed. The use of AI in recruitment is already common, and is growing fast (Deloitte, 2019).

Text processing methods can be used to do a variety of things such as automatically extract qualifications, skills, expertise; match candidates to jobs; and assess candidates for a particular job. Neural networks can, for example, take a particular phrase from a CV, analyse the context for the phrase based on surrounding words and determine whether the phrase is a skill defined for the particular job.

Natural language processing tools can be used for text messaging based interviews to capture meaningful information using semantic parsing named entity recognition and multiple intent classification (for example Mya, a conversational recruitment tool). Machine learning algorithms can learn continuously from millions of interactions and constantly develop and improve accuracy.

AI systems can also be used to classify non-verbal behaviours in job interviews. Such AI systems can identify facial expressions, speech signals and visual cues such as eye gaze. This technology is still at the early stages of trying to interpret non-verbal behaviours in a meaningful way and there is quite a high degree of scepticism. While no system appears yet to be operating without human oversight, there does appear to be increasing use of these tools to help with applicant screening to create rankings of applicants. It is difficult to find real examples of these AI tools online, whereas there are a lot of demonstrations available. Some thought this may be a sign the technology is still developing and is not yet very good.

Colin Gavaghan's presentation took a step further back along the recruitment timeline. The "hiring funnel" begins at the "sourcing" stage, with determinations as to who even learns about vacancies through targeted advertising. AI is already playing a major part in this upstream part of the process (Upturn, 2018). While this offers employers advantages in terms of reduced time-to-hire and hiring-related costs, concerns around bias, accuracy and such like begin at an earlier stage than is sometimes realised. This decisions may also prove to be highly elusive regulatory targets, as those screened out by such tools will have no reason to know that they have been "pre-deselected" in this way.

AI and regulation

A key theme (which was echoed in other sessions) was that these AI tools often led to reflecting how well humans do and how human decision-making is scrutinised. For example, participants asked how well do humans perform in selection and recruitment processes and raised concerns about existing gender, race, cultural and other biases. Another theme was

that AI tools which could assist humans to make better decisions should be welcomed, particularly if these were able to counter some effects of human bias, but such tools needed to be used in circumstances where the risks were well understood and properly mitigated.

In some areas AI tools would not be useful in some areas of the professions (such as for highly specialised professional skills) or in smaller workplaces. The more common use of AI tools is currently in large companies because these contexts offer economies of scale. This raised the question of how many large corporates there are in New Zealand and how many were large enough to use these kinds of tools. Utility might also depend on how many people an agency was wanting to appoint or how many might apply (for example one existing use is to assist with preliminary screening of scholarship applications). There was also the possibility that large global corporates with a small presence in New Zealand might use AI tools when selecting New Zealand employees.

Participants wondered whether the development of AI tools was being done primarily from an employer perspective, rather than from an employee perspective or in a citizen centric way. With a different approach, for example, could AI tools be used to ensure that the jobs advertised were well designed? There were concerns about the number of candidates who ruled erroneously themselves out of job opportunities because they felt they were not qualified or the job description reflected poorly on the agency. Might better job design, with assistance from AI tools, help widen the pool of candidates and improve corporate management practice?

Textual AI systems come with risks and may be able to be circumvented, for example, by including key words such as 'Oxford' or 'Cambridge' in CV in white text which can be picked up by AI textual submission. Such AI systems exist overseas, but it was unclear if companies in New Zealand were using these or developing their own. In the case of overseas systems, key questions to ask include: will the data set be relevant in the New Zealand context? How much subjectivity might there be in choosing phrases relevant for skills in New Zealand and how might developers know whether or not these words, which are included in textual selection systems, signal a particular skill. How might employers using such tools deal with concerns about subjectivity? One option is to have multiple annotators and then cross correlate these so that if there was a higher or lower degree of agreement this might, or might not, signal bias.

On the one hand AI tools might improve fairness because there is an opportunity to confirm a classification that has been made through human oversight - if a classification does not make sense then, just like a human, these can be further examined. One suggestion was that companies developing these tools strengthen their transparency, for example, by publishing the classifiers they use or considering reporting or other forms of public disclosure (for example, a requirement to disclose that a candidate is not talking with a human).

Participants discussed existing research on recruitment bias in AI tools and wondered whether hiring with the assistance of an AI system might incentivise employers to use more objective criteria. If so, would this be fairer and therefore better overall? What might be the implications of replacing one imperfect system with another? Would employers overall be

recruiting more people and increasing diversity or would these tools only amplify and further embed existing biased based on poor quality data sets?

Participants wondered how AI tools would comply with general legal obligations, for example, human rights and non-discrimination. The issue of compliance with more specific obligations was also raised, such as the Scandinavian laws governing gender diversity and other prescribed requirements for diversity targets. Could an AI hiring or recruitment model be trained to find bias in a model? Human bias was also discussed and participants noted how some corporates are looking at male and female encoded language in job descriptions to try and eliminate gender bias. On the other hand, if an employer wants to increase diversity how might that be achieved with the use of these tools? This led to a wider discussion of whether there is social agreement on the ideal state of participation and representation of society in employment – is the goal a certain percentage of diversity, how does progress get measured over time and how does this fit into discussions taking place in the professions in New Zealand?

Finally, participants considered the use of AI tools raised more fundamental issues about current hiring and recruitment practice, for example:

- can the tools be used to identify skills gaps? For example, instead of filtering out candidates, can some of these tools be used to find more workers in those parts of the workforce where there are gaps?
- what is the real purpose of recruitment?
- cover letters and CVs represent “old technology” and outdated practices - should we be simply mimicking these existing HR processes or using new AI technologies to create entirely new ones?
- why is there an obsession with finding “the perfect candidate”?
- is a fast decision to select an imperfect candidate preferable to extended hiring processes that search for the elusive perfect candidate?
- what, in practice, is the correlation between a person’s CVs and their actual job performance and how is this taken account of in candidate selection and testing?
- how might applicants circumvent these tools?
- how many people will really be affected? Some jobs are never advertised, some filled through direct shoulder tapping and so on.

AI and employment law

This session focused on the employment law challenges posed by AI, including workforce management, redundancy and replacement. Gordon Anderson (Victoria University of Wellington) and Avalon Kent (Council of Trade Unions) provided context for this discussion.

Gordon Anderson outlined an immediate challenge was the ever increasing use of AI to actively manage workers and the related employment law issues this raises. This was seen as a more pressing and significant challenge than risks posed by wholesale redundancy and iterative replacement. Rogers (2019) for example, argues that AI related automation is not a major threat to workers today nor in the near future. Rather, new information technologies are being used to give employers a new power advantage over workers and these are able

to be implemented under existing employment laws. This monitoring includes a wide variety of activities such as using algorithms to monitor, direct, or schedule workers, thereby reducing workers' wages or their work-related or "on the job" autonomy. Participants heard that companies are also using new technologies to "fissure" employment, which occurs where work tasks or processes are outsourced and related legal duties towards workers are denied, while at the same time the employer begins to more closely monitor workers' performance.

Standardisation of cross-border human resource management (e-HRM) practices is exacerbating this, with some commentators (for example, Dörrenbächer et al), concerned that:

1. The expansion of standardisation into strategic HRM, through human capital management systems.
2. Such systems not only may have a negative impact on employees but may largely disempower national level HR departments.
3. The dangers of standardisation are identified as not only their increasing reach but their objective by turning HRM "into a value-creating will corporate function," (which is not likely to be positive for employees increasingly micro-managed at the individual level).
4. Such measures can and have already been used by management to discipline workers, weaken their overall bargaining position, and resist calls for higher wages.
5. On the level of individual jobs and work tasks, we also see an increased risk of a loss in decision-making autonomy, especially when strict 'lean production' regimes are enforced and digital technologies take over controlling tasks which until now had been performed by specialised employees.
6. This is often bolstered by the implementation of sophisticated software systems which are used for tightening technological control of performance measures for individual employees, teams, and whole production and administrative processes.
7. e-HRM lays bare the growing risks of intensified surveillance of employee behaviour at work and beyond, which includes their spare time.

Implications for New Zealand employment law included loss of local or national corporate autonomy in employment relations, weakened position of workers in the employment relationship (with less autonomy and more monitoring) and tuning AI systems to minimise statutory benefits such as holiday pay.

Many of these changes to employment were happening under existing law and one question was whether it is too late to regulate – increased monitoring is already taking place. In employment law, AI adoption and implementation by management falls within "managerial prerogative" and the obligations of employees to obey reasonable orders of their employer. In New Zealand, the judiciary have been very reluctant to interfere with managerial prerogative and, if current practice continues then addressing this regulatory deficit may be problematic.

An assessment of the employment related legal implications of AI was also needed to take into account the wider context of New Zealand employment law over the last 20 years. This

includes, for example, the privatisation of employment law and the significant democratic gap which has opened up in the changed situation of trade unions and the lack of an organised voice for workers. In addition, New Zealand does not have a large number of multi-national companies and smaller employers do not yet use cloud systems at scale. This may afford New Zealand some protection from the full impact of AI tools on employment and provide a window of time to consider regulatory measures more deeply. One option suggested was to provide a stronger statutory basis for consideration of the rights of workers to dignity, privacy and to a family life within the statutory duty of “good faith” employment relations.

Avalon Kent posed the question: what is the role of trade unions in a discussion about AI and employment? Answering this, some thought, needed a socio-political analysis and to consider the impact of AI and work in the context of power relations in the employment relationship (which is acknowledged in section 3 of the Employment Relations Act). Rather than talk about the future of work, the discussion should be humanised and focus on the future of workers. Seen in this way, it was thought, there may be opportunities to better understand who drives change, why and what forms of accountability there might be.

In general the trade union movement does not have an alarmist view of AI. Nor is a binary view of whether AI as “good” or “bad” helpful. Instead, it is preferable to see AI and both an opportunity and a potential threat. For example, if AI related processes are able to reduce dangerous work for humans, that is a good development, but if these create more oversight of employees they may be unfair. It was important to consider the impact on different sectors as some may be more prone to the effects of new uses of AI than others.

In relation to redundancy and replacement concerns, OECD research shows that New Zealand is at less risk of automation than other OECD countries and that change is more likely at the task rather than job level. A more likely scenario is that AI will supplement, augment or support work. However there may be some unintended consequences. For example, job polarisation may increase between those who have the skills to move between jobs and those who do not. There will be a need to ensure workers are educated about AI and area “AI literate” to help prepare them for working alongside new technologies.

Australian research, particularly in the steel industry, indicates that AI doesn’t necessarily lead to more exploitative practices, but that this depends on *how* rather than *whether* AI is deployed. A key part of this example was the involvement of unions in implementing new practices. In adapting to this new context trade unions have supported emerging ethical principles such as the *Top Ten Principles for Ethical Artificial Intelligence* (UNI Global Union, Switzerland). Other work related regulatory areas for focus might be health and safety, redeployment policies and wider social security policies such as the Universal Basic Income. Finally, AI deployment could generate resentment and this could be exacerbated by a failure of regulation to keep pace. The New Zealand Future of Work Forum was an important initiative in this context.

Participants noted sector losses a very important area to watch and thought the Productivity Commission’s use of scenarios in its discussion document on the future of work was helpful. Some work could be done to look for triggering signals which AI experts might agree should

alert a prompt response if these occur. Participants considered the use of machine learning tools to assist work related accident compensation research would be useful , although care would be needed as reported accidents do not measure workplace harm and or near misses.

Unexpected algorithmic collusion or unworkable output results that emerged where humans were not involved also needed to be monitored. Participants shared two such examples: (1) A supermarket in the United States of America which used AI tools for scheduling and then scaled back the use as it found a wide range of unexpected scenarios the tool couldn't cope with (for example distinguishing between coconuts and cupcakes) and research findings that use was feasible for only certain predictable rule based tasks; and (2) Use of an AI scheduling tool in a hospital which assigned minutes per patient to nurses but had to be scaled back as it was not effective for the nurses in practice and it was preferable for nurses to determine how long to spend with patients.

Participants distinguished between management of work and management of people. The distinction is fundamental in employment law and in the New Zealand context of regulation of "employment relations". New Zealand law regulates the employment relationship, which is one where the parties may not have congruent interests (much like in other relationships). Employment relationship legal concepts are founded on *human to human* and to that extent AI cannot itself be an employer. Employment law also concerns justice in a workplace which cannot be tabulated into a checklist of the kind which AI tools are designed to use. There is extensive New Zealand case law supporting the need for an holistic and contextual application of the law in any given factual situation. For this reason, it was doubted that AI could achieve employment related justice on its own.

There were concerns about assumption that employee monitoring was acceptable and about what duties there were on employers who had information (for example if the information revealed actual or risks of health problems for employees, then the employer's health and safety obligations would arise to identify, assess and take all reasonable steps to mitigate the risk. Overly intrusive collection was also a serious concern, for example, orderlies in hospitals having to wear a monitor that tracked all their movements, including personal ones such as going to the toilet. Another example cited was Ovio, a US company, where there were so few women employees that it would be possible to identify each one, so the women aggregated the data to avoid individual identification.

Insofar as the technologies were being used they may be options for designing specific rules in some areas. One option suggested, for example, in relation into use of employee information by AI in hiring and recruitment or employee monitoring was a new, specific code of practice under the Privacy Act 1993. Another was the creation of data trusts, managed by trustees from unions and employers, through which employees share information that can be used. Another was to introduce more employee participation in existing corporate structures, for example, to have union representation at Board level or a representative supervisory board, above the company board.

More fundamentally, participants asked whether we should regulate for technology or adapt technology to the regulatory framework? If facial recognition use picks up that a

person is anxious and draws adverse inferences and conclusions, for example, there could be discrimination minefields for employers. In this scenario is it preferable to change discrimination law or to leave the current regulatory burden of compliance with non-discrimination laws on employers.

Other human rights concerns were noted including that the right to join a union is not effective enough in New Zealand at the moment and that has implications for how effective other human rights related interventions might be. One E Tu union initiative was to negotiate a number of ethical issues with an employer including the minimum wage for workers using new technologies and this was done despite there being no regulatory requirements. In other cases, many workers self organise on social media platforms (for example, the are New Zealand Facebook groups of 15,000 truck drivers and 9,500 Philippines farmers) as a way to deal with the power imbalance in employment relationships that do not have union representation. For example, employees might collaborate to turn AI tools off.

The power imbalance in employment relations was considered complex, with many parts. Recent case law involving use of algorithms in employment contexts was discussed including *Houston Federation of School Teachers v Houston Independent School District* 251 F.Supp.3d 1168 (2017) (algorithmic tools in public school teacher evaluation). A key difficulty was access to proprietary information about how algorithmic tools worked, information which is commercially protected under US intellectual property laws.

Employment and the gig economy

Rakesh Mistry (Straker Translations) and Elizabeth George (University of Auckland) provided both a practical example and academic critique of employment and the gig economy in New Zealand.

Rakesh outlined the context for development and use of algorithmic tools for translation services. Demand for translation services is high and AI tools can produce a draft translation which is reviewed by a human. With over 10,000 translators available globally, the translators bid for projects and are assessed for their skills to help with project matching. Staff have independence, flexibility and while there is no set pay rate benefits included that work was ad hoc, freelance and short term. Translators are evaluated (although workers might not know their scores) and complaints are subject to peer review.

Elizabeth George, building on the theme of employment relations and power, considered the myth of choice, flexibility and mobility in the gig economy and the direct and inadvertent parts which employees may play in driving this myth. Concepts of non-standard work and non-standard workers (those in temporary, mobile or part-time work) were considered relatively unimportant for a business. Instead, work on these arrangements would better proceed by understanding that worker motivations are diverse, that greater choice means greater satisfaction and career outcomes are often limited. Organisational considerations relate primarily to flexibility (for example for numerical and functional for service or product demand).

New regulatory proposals have emerged, for example, in the US non-standard workers cannot be employed to do the same job standard workers. Some research shows that standard workers don't like working with non-standard workers (although the opposite is not true).

New research is being carried out on how non-standard workers in the gig economy are affected by isolation and loneliness as well as research on the movement of workers from non-standard to standard work. One study found that mobility is extremely low, that most workers who were non-standard wanted to be standard workers, that non-standard work is only a stepping stone into standard work for a small group. In relation to the benefit of increased flexibility from gig economy work, research showed that women and parents are the most likely to switch from standard to non-standard work but when they do so, they seldom switch back. The results raised more questions such as whether the "boundaryless career is bounded by work status" and while non-standard jobs are not always bad, it is possible to promote mobility and empower people in non-standard work?

Some research does not affirm the view that there are significant potential cost savings for use of non-standard workers, with in-house research most often not being published if it reveals adverse results. A particular area of concern was the tension that arises when permanent standard and temporary non-standard workers sit alongside each other doing the same job.

Participants saw some risks that there might be a backlash against gig economy work because of bad behaviour in some companies. More fundamentally, were new forms of gig economy work the beginning of the end of standard work – would this be a 'slow earthquake' which saw movement away from standard work over a period until it finally moved completely? What were the implications for families of non-standard work arrangements? New AI employment tools could be a helping hand or a threat depending on whether these helped humans to do a task or took over the task from humans.

Ethical and regulatory challenges

The ethical and regulatory challenges relating to AI and employment were assisted by presentations from Matt Boyd (Adapt Research), Toby Gee (Barrister) and Richard Wallace (Parliamentary Counsel Office).

Matt Boyd discussed the ethical challenges with reference to the scenario of contract cleaning work in the Uber business model and analysis of algorithmic based work allocation (Moore, Upchurch and Whittaker *Humans and machines at work: monitoring, surveillance and automation in contemporary capitalism* (2017); Raso, Hilligos et al *AI and human rights* (2018) and Susskind's *Algorithmic Injustice (Future Politics)*). A key concern was that neutral algorithms can operate to maintain and embed injustice in an unjust world.

There may be new rights that flow from the new technologies and these will also need ethical consideration including at technical levels (the IEEE has developed *Ethically aligned design* principles) which need to be seen in the context of major ethical theories. A deeper ethical analysis is needed of why is surveillance wrong to better see issues of the power

imbalance (equivalence between employee and employer), the inability to “shirk” or take micro-breaks and the fears that data will be misinterpreted.

These raised ethical questions such as: should you be allowed to forgo your privacy; is there a collective privacy that is greater than the sum of individuals’ choices; can you opt out; can customer satisfaction balance out harms to employees; and how can we make the judgements of customers and employees commensurable?

Toby Gee outlined the traditional approach to employer responsibility for employee use of technology and the employer’s obligation was the duty to use reasonable steps to avoid injury (the United Kingdom case of *Stark* a case in which an employee was injured while using a bicycle). However, the Enterprise and Regulatory Reform Act 2013 changed this general approach so that a breach of a regulation does not result in liability unless the employee can show negligence.

A similar approach applies to algorithmic use which would lead to the result that the producer of the algorithm is liable in the same way as for any other products. New Zealand has moved to allowing software and other intangible property to count as goods, rather than services, which would support this approach.

In relation to regulatory implications, reference was made to the traditional common law approach and to ask in the first instance: does algorithm use make any difference in our contractual or other legal duties (like hiring)? One way is that algorithms might be opaque in a way that other processes are not, which may suggest a different approach is needed. From an employment law perspective, it might be thought that an employer will fail to demonstrate good faith employment relations if he or she is using algorithms that cannot be explained to employees. A further regulatory question is: if an algorithm is found to be biased, who is liable? Should general duty apply so that an employer only has to do what they reasonably can OR should there be a strict duty of some kind? While this is a policy decision, it was noted that New Zealand does have consumer protection laws which have strict liability for harms caused by faulty goods. The new GDPR rights (such as the right to a decision that is not solely algorithmic) were also noted.

Richard Wallace reflected on the development of statute law, noting one view was that regulation should be aimed against harms not technologies and that an AI specific law might not be needed. On the other hand, law makers do frequently regulate human reasoning processes, so why not artificial reasoning?

The use of technologies to assist legal compliance, for example, through the production of legislation in machine-readable as well as human readable form (rules as code), was also discussed. This approach might work very well in areas of the law that are highly predictable. A number of barriers to automating processes were revealed such as legislative references to an action being done by a human, the different definitions of “child” in different statutes and so on.

In the New Zealand statute book, nine laws currently relate to AI and automated decision-making specifically – see for example section 296 of the Customs Act. There are a variety of

other references to algorithms, which are not separately defined. One regulatory option suggested was to have one piece of legislation for digital matters and update that as technology evolves, rather than amend various parts of the statute book every year.

AI in the legal profession

The use of AI in the legal profession is still new, but is growing. Presentations by Jean Yang (McCarthy Finch), Matt Bartlett and Geoff Simpson (CitizenAI) and Mary Ollivier (New Zealand Law Society) set out the current context in New Zealand.

Jean Yang noted that drivers for new services include the high rates of people with unmet legal need and the very low number who will engage a lawyer to help (88% of people with a legal problem will not engage a lawyer) and the overwhelming barrier is cost. Legal business models are changing with the decline in the billable hour, resistance to fees for manual tasks such as photocopying, and the increasing view of legal advice as a roadblock to business rather than an enabler. In this environment, there is opportunity to productise advice rather than offer legal services and to offer set fees. Other drivers for change in the legal profession more generally are the demands for culture and career path change and the desire to encourage experimentation and democratise access to legal information. The recently established New Zealand Law Workers Union was noted and participants wondered if the more positive views about the uses of AI tools might assist. For example, using these tools to help measure actual working hours as opposed to billable hours in order to assess whether law firms were paying the minimum wage.

Most products are working with narrow forms of AI, such as natural language processing, which can be applied well and are delivering most impact. Development is data hungry which can create a barrier (although there may be options for data sharing to be explored), mastering nuance is different and made more difficult by the varieties of legal reasoning and explainability issues. However, new processes can be very effective and this is a growing field in New Zealand although still in the research and development phase.

The impacts on employment in the legal profession need wider discussion such as how to train juniors, where juniors fit into a law firm business model and the options they have for more meaningful work and skills development. In the future, there will be a need for more lawyers who speak the language of technology and who can appreciate and assist with ethical questions. More generally, new ethical questions may be asked of practitioners, for example, should there be a professional development obligation in relation to new technologies?

Questions may arise in the area of “machine led human governance” that pose new legal difficulties if lawyers are working with black box AI tools. For example, how do the corporate governors lawyers check transactions or how do data scientists checking the technology; is a lawyer in the loop really “monitoring” (decision atrophy, attention issues)?

In New Zealand new forms of legal information service have already developed, such as CitizenAI. Based on values of social participation and active citizenship, para-legal chatbots have been created for particular areas of law, such as RentBot. These use Google dialog flow to create tools for simple queries. Complex cases are handled by referral to a human.

Mary Ollivier noted that the development of new AI tools is occurring in a fairly open regulatory profession environment. There are many new back office legal technologies emerging to help make legal practice more efficient. But client facing (as opposed to back office) legal tech is still small in New Zealand for many reason (isolation, separate jurisdictional issues, small market, large international firms not here). There is no “Xero” for the legal profession in New Zealand yet. There is also a current prohibition on NZ lawyers sharing income with non-lawyers. But the reserved areas of law for lawyers in NZ are actually very narrow and the profession growing by 500 lawyers per year. It is unclear if there are gig economy law jobs coming.

As a regulator, the New Zealand Law Society takes a positive approach, working those firms who are developing new technologies to support compliance within existing rules, although there are no drones filing court documents yet. The regulatory approach has been that where the service is a regulated one (bearing in mind the small range of services that are regulated – legal advice and preparation of court proceedings, conveyancing and property relationship law), there must be a real person behind it and one who is authorised to practice on their own account. That person has all the normal obligations to supervise its operation and is responsible for it – a professional competence approach.

Other jurisdictions such as Australia, England and the United States are taking a closer look at the legal profession and new technologies. A recent independent review of the regulatory framework in England and Wales by Prof Steven Mason, for example, found that about 80% of legal work that is not regulated and made a variety of recommendations, including relating to law and technology. Also in England and Wales was Alison Hook’s report (Hook Taganza) for the Legal Services Board, “The Use and Regulation of Technology in the Legal Sector beyond England and Wales”. A recent Chief Justices conference in USA that urged States stay in touch with technological developments, making a variety of recommendation for doing so including developing technology strategies; setting up an international group on ethical issues; pooling resources; and building their own knowledge of technologies.

Lord Thomas of Cwmgiedd noted that the distinction between using information technologies for legal processes and using these for giving legal advice was a useful one. The high cost of legal services was creating significant difficulties and, in the United Kingdom, government funded legal aid was for a narrow range of cases such as housing, social services and family disputes involving children. In New Zealand regulation may not be a problem, but rather how to get legal advice to ordinary people, which is where p AI has real potential. Finally, when considering humans in the loop, it was suggested that the purpose of oversight needed more assessment: how high are the stake and how vulnerable is the person receiving the advice? In a regulatory context, the aim of professional regulation was to protect consumers, rather than employees.

AI and health care

Participants considered current research and use of AI and health care with assistance from presentations by Pieta Brown (Orion Health), Elizabeth Broadbent (Auckland University) and Angela Ballantyne (University of Otago).

Pieta Brown discussed how machine learning research and product teams are working in New Zealand contexts to develop and deliver tailored client solutions in healthcare. These tools are working to deliver capabilities we often associate with human intelligence such as sensing, reasoning and acting to support improved health outcomes and efficiency. Examples include vital sign monitoring, predicting surgical outcomes and deep learning from medical records. The foundation for this work is health data with the result that health care is increasingly becoming a data science.

Some issues are emerging in this area including how to get access to data, the complexity of the health ecosystem, risk aversion, issues with data quality and bias and issues with providing model interpretability (in a way that supports clinical action). Other issues related to use of new tools by medical professionals who are time pressured and naturally resistant to adding more screens or clicks to their workflows and some real issues with practical use as opposed to research or study have emerged (over-trusting of automated support systems, hyper-vigilance in study contexts, bias and lack of transparency).

Nonetheless, new questions are emerging as the technology develops quickly and even as there are skills shortages in some areas. For example, if diagnostic imaging tools improve, should we stop training radiologists or other such specialists? The answer at this stage was no, since such specialists do a lot more than detecting patterns in pictures but the new technologies might propel people away from these occupations which, perversely, could lead to skills shortages in some sunset areas of expertise.

Elizabeth Broadbent discussed current research on human-robot interaction in health contexts, including studies on expectations and attitudes towards robots. Research shows that some people do achieve companionship with robots, have generally positive perceptions about what robots might do for them and about benefits of using them (accuracy, savings in costs and time of travel to medical services and so on). However, people also have fears such as loss of jobs, interferences with privacy, of accidents and worries about the effectiveness of the technology. Research into use of care robots in rest homes showed that managers were very positive about the prospects of their use as there is significant unmet demand for care workers.

In a rest home context, a care robot was placed in the lounge supplying games, music and supported other care activities such as taking blood pressure and giving reminders to take medicine or other health related actions. The research found that care robots, such as Paro (which looks like a baby seal), reduced loneliness for some patients and reduced agitation in patients with dementia. Robots in the home did help with loneliness and might, in some cases, have encouraged visits by extended family, such as grandchildren, who wanted to see the robot.

Angela Ballantyne provided an overview of ethical issues for AI in medicine. In doing so, there was a need to consider 'relative' risks and benefits (so there is no double-standard in comparing to existing practice). A strong emphasis was put on examining the reliability and accuracy of AI tools rather than their explainability as there are many technologies that practitioners use which they cannot explain, but which they are satisfied are reliable and

accurate. In this regard, the profession is aiming for that “sweet spot” at the intersection between law, ethics, and social license.

Benefits of AI tools include clinical utility (medical error is the third highest cause of death in the United States of America) and the potential for cost savings given the massive increase in healthcare costs which in some countries is outpacing GDP). AI tools can far exceed human accuracy in imaging diagnostics for skin lesions, pathology slides, chest radiographs and in detecting retinopathy. However, human and technology interaction is still needed, and clinicians work with the results to understand what data really means in a particular case.

Ethical issues of transparency or explainability need to be considered carefully. Doctors can give explanations of their reasoning but these are often inaccurate and may be incomplete. Doctors are already influenced by many biases including availability of particular treatments. More generally, much of medicine (e.g. epidemiology) rests on correlation, not causation, for explanation.

In relation to bias, significant risks arise where algorithms trained on datasets that reflect inequalities and conceptions of “normal”. Bias in health AI may also be affected by being trained on biased historical health decisions.

Access to health data and consent to use of health information is another significant issue. The investigation of the UK Information Commissioner into the case of the Royal Free Hospital in London granting Google access to 1.6 million patient records found that:

- consent should have been sought for Deepmind to have access to the health information
- the NHS should have been more transparent given public expectations (while consent is often not possible in health contexts the public in this case should be informed).
- the NHS shouldn't have *gifted* the data to Deepmind.

Reflecting on the New Zealand context two areas were relevant: the Health and Disability Code of Patient Rights (Rule 6 – what information would a reasonable patient need to understand the technology) and the Health Information Privacy Code (Rules 10 and 11(c) – is the disclosure for a purpose for which the information was obtained). While these provide a good basis from which to develop ethical practice in use of AI tools, there was also a need to ensure public support and trust as the public have particular concerns about commercial use of their data, particularly sensitive health information.

Participants discussed concerns about commercial access to data and whether this was really a concern about data broking. Research shows that public concerns about commercial use include harm to them or their family, exploitation, data breach. In addition, as a matter of principle, the research showed a strong view that public data should not be exploited for private commercial gain. Concerns about Maori data sovereignty were also being raised.

Participants discussed the implications of use of AI tools for unmet health needs alongside unmet legal and other needs and how this increasing use might assist in building public trust. In this context, explainability could help improve trust at least in the early stages of the technologies development, but rigour is essential and some standards for explainability might be needed so that comparisons are meaningful. For example, what level of explainability is needed may vary given that people want different explanations for different reasons.

In relation to ethics of care robots, participants wondered how anthropomorphism enables or restricts human interaction with robots. At this stage, the main ethical consideration was that the person perceives being treated with empathy, not whether the robot actually does. People interact differently – some are more open with a robot than with a real person while some people did not want to interact with a care robot at all.

One interesting observation is that AI tools are clearly moving into the emotional task space as well as the mundane task space. This seems at odds with the general view that such tools would take over more mundane, repetitive and automatic tasks that would lack emotional content. However, some thought that it wasn't possible to predetermine which space AI would be used for as human also influence that by how they interact. How and where to deploy was an open question as the tools were effective in both spaces. For example, such emotional care by AI tools might assist health professionals who can suffer burn out through dealing with high levels of mundane emotional labour.

On the other hand, might such tools change how humans interact with other human health professionals? For example, if we expect non-human systems will take care of certain tasks, might we treat humans less or more humanely? Concerns about research on social isolation were raised and the implications of trying to solve this through more interaction with non-human AI tools (and related interaction such as fixing it, updating it and so on).

A related insight was that it is one thing to gauge the impact of AI tools in health contexts *now*, and another thing to imagine forward to a day when AI tools are a lot more prevalent in many different areas of life. AI tools can appear to have a largely positive effect if they are used in moderation, but the effect might be quite different if they become the dominant mode of interaction. For instance, a few cuddly AI toys in a retirement home may have a fairly unambiguously positive effect. But a retirement home where many of the people's needs are fulfilled by robots is another matter altogether. This seems to be exacerbating the current trend towards social isolation, loneliness, and associated psychological problems in a dangerous way.

Digital assistants and counsellors

Participants discussed the development and deployment of AI digital assistants and counsellors with assistance from presentations by Jim Warren (University of Auckland), Gareth Cronin, Lena Waizenegger (Auckland University of Technology) and Hazel Bradshaw (Department of Internal Affairs).

Jim Warren provided an overview of digital assistants relating to mental health and young people in particular, including:

- eCHAT (case finding and help assessment tool)
 - Designed to be used in the waiting room
 - Module structure is: screening, discovery, help, next module
 - Acceptable tool in general practice
- SPARX (gamified self-help,)
 - More available to young people and less stigma
 - Assessed as “not worse” than face-to-face.
- HABITs (behavioural intervention tech)
 - Modules: remind, relax, reenergise, rethink, resolve and relax.
- HeadStrong (chatbot)
 - Based on cognitive behavioural therapy and positive psychology
 - Delivered through FB Messenger

These systems can be used in a variety of contexts and use dialogue flow techniques. Chatbot exercises with users are generally brief and do not deal with their individual cases (because of the limitations of the AI). But they are useful for reinforcing and reminding people of tasks such as taking a break, relaxation and so on. Such systems are not trained on counselling session data and at this stage the results are poor as it is difficult to allow a such training system to ‘play itself’ as happens when training for a game, such as chess. There are some accredited health information sites (such as Health on the Net Foundation), which might be able to move into accreditation of these kinds of digital assistants.

Gareth Cronin provided an overview of developing commercial and financial services chat bots and conversational AI tools and moving on to make systems for building such bots. These systems are now used to make chatbots in areas such as health, media, ACC, finance, retail and telecommunications. The aim is to deliver better customer experience by using analytics from real conversations to give developers and business owners much better information about what customers think and want. Research shows that chatbots are only most effective when this is the only channel for customer communication. One practical challenge for designers was how to end a conversation with a chat bot.

Ethical questions include whether a user should be informed that their data is being collected and may be used for purposes other than solving the user’s problem. A solution to this was the use a statement modelled on the GDPR requirement for transparency.

Lena Waizenegger presented research on the impacts on human employees from conversation agents. Augmented intelligence systems are used to provide short term customer support and content curation or longer term services such as coaches and personal assistants. As with AI and health services, these tools step into the emotional space, offering emotional support for human activities.

Participants considered care was needed not to conflate different concepts. For example, to consider the complement of technical routine tasks is, emotional or empathetic reasoning. Humans also provide a wide variety of other emotional tasks such as critical thinking,

negotiation, creative thinking, complex communication, ethical / political discourse. These tasks require a much wider breadth of knowledge and semantic understanding.

As with other AI tools in other areas, the results from digital assistants are mixed. Some people get better personalisation talking to bots as they are more open and honest in the information they supply. In other cases there may be less task variety but more task significance and therefore a need for higher accountability.

Hazel Bradshaw outlined the context for work on AI tools in the programme of the Government Chief Digital Officer. New Zealand does not yet have an AI strategy and has a 2015 ICT strategy. A key consideration for future work is the emerging technology landscape. Looking ahead 20 years this emerging technology includes:

- AI tools and applications
- spatial computing
- encryption
- IOT
- Big data
- Robotics and autonomous systems
- quantum computing

There are some interesting possibilities for the use of AI tools for policy development. For example, following the Christchurch terrorist attacks in March 2019, Hazel has developed a virtual reality hate speech simulation. The purpose is to demonstrate how terrorist and extreme content can influence people and measure the effectiveness of interventions in a virtual environment.

Participants reflected on the issue of AI tools that appear to empathise and whether this amounts to impersonating a human. Humans can treat other human staff in customer service centres very badly – might it be preferable to reduce some of the emotional labour on those staff? Other questions were: is a user knowing that the digital assistant is not a human is better for that AI tool and whether interaction might affect how humans express emotions (we already use emoticons for example).

Participants wondered whether there could be a standards authority for chat bots as another route for building consumer trust. For example, if a robot can exhibit professional behaviours should it be required to meet the same professional standards as practitioners in the relevant field (health, legal, education, journalism or corporate management fields? Could there be a limit on how many times the chatbot might try to change a customer's mind and then stop?

Conclusion

Several strong themes emerged including the clear potential benefits for AI in the professions. Development and deployment is underway and there is scope for more interdisciplinary and multi-stakeholder collaboration to share insights and experiences as the

field develops. Overall a principled and pragmatic approach emerged, with participants welcoming prospects of new technologies, but very aware of the limitations and risks.

The idea of a large 'unmet need' seemed to be a common theme across professions. There was concern over levels of unmet legal need, for example, the many barriers to accessing professional legal services and the challenges facing the legal profession to change its culture including how it trained junior lawyers. The use of AI tools to improve access to justice (rather than simply to make law firms more efficient) was seen as a significant benefit, particularly given the large scope for providing unregulated legal services.

Concerns over unmet needs were echoed in the health profession as were the views of the potential benefits of using AI tools, including care robots to address these unmet needs. The use of these tools raised significant ethical issues, including the range of emotional work that can be performed by robots in light of how people engage with them.

Participants considered the implications for AI and employment and jobs were not likely to be tectonic, but that change could happen in a 'slowquake' fashion. For example there could be a gradual loss of standard work. There were concerns about the implications for the employer – employee relationship. AI tools also offered as new opportunities that might help human workers, employers and other professionals to change their perspectives. For example, could AI help reveal bias? Might it be possible to develop the concept of "management prerogative", or to change the 'employee as property' model of employment relations. AI tools might help employers to think differently about work and job design, for example, in identifying classifiers for core skills and competencies in job descriptions. More fundamentally, employers may be challenged to change candidate selection techniques and debunk the "myth of the perfect candidate".

A cautionary note was sounded for workers over the prospect of AI tools might not provide more work flexibility, cost savings may not eventuate and the risk of public discourse being informed by a narrow, rather than wide, set of research results. Involvement of workers and trade unions in development and deployment of AI tools and systems was encouraged based on some overseas initiatives.

In many areas, AI systems are not in any danger of replacing people - because there are already no resources to employ people to do all the work that is needed. If AI systems address an unmet but vital need, there is arguably less concern about AI replacing human workers - and also perhaps less urgency in the call for their oversight. On the other hand, commercial perspectives may shape the direction of future investment depending on responses to the question of how much money there is in addressing unmet needs.

A diverse range of ethical issues emerged particularly around the plethora of newly developed ethical principles to assist in diverse sectors. While some these ethical issues relate to well known areas (such as bias and control) some new areas emerged in relation to the emotional work that digital assistant AI tools do.

Development and deployment of AI tools is taking place within the existing regulatory schemes for the legal and health professions and under existing employment law. Legal

changes were occurring as needed and, while there were no significant gaps, participants did see some cracks and potential for problems as the use of AI for more tasks increases.

Ideas for regulatory initiatives included:

- Developing legislative models where AI is used in automated decision-making
- A code of practice under the Privacy Act
- Whether there is a role for professional regulators to monitor how AI tools are being used and the effects on their professions
- Whether there is a role for professional regulators to monitor how AI tools are meeting unmet legal, health or other needs or to improve health outcomes or access to justice
- Whether standards for digital assistants could be introduced.
- Whether discussion of new standards should focus more on reliability than explainability of AI tools.

Appendix A: Roundtable Programme and Participants

Thursday 26 September 2019

8.30 - 9am Registration, welcome tea and coffee on arrival

9am-10.30am

Welcome and orientation [9.00 – 9.15]

Uses of AI in hiring and recruitment

Alistair Knott *University of Otago* [9.15 - 9.35]

James Maclaurin and Colin Gavaghan *University of Otago* [9.35 – 9.55]

Discussion [9.55 – 10.30]

10.30 – 11.00 Morning Tea

11.00 – 12.00 pm

Laws relevant to AI-related redundancy and replacement

Gordon Anderson *Victoria University of Wellington* [11.00 - 11.20]

Avalon Kent *Combined Trade Unions* [11.20 – 11.40]

Discussion [11.40 – 1 pm]

1 pm – 2 pm Lunch (provided) [1 pm – 2 pm]

2 pm – 3.00 pm

Employment and the gig economy

Rakesh Mistry *Straker Translations* [2.00 – 2.20]

Elizabeth George *Auckland University* [2.20 – 2.40]

Discussion [2.40 – 3.00]

3 pm – 3.30pm Afternoon Tea

3.30 – 5.00 pm

Ethical and regulatory challenges

Matt Boyd *Adapt Research Limited* [3.30 – 3.50]

Toby Gee *Barrister* [3.50 – 4.10]

Richard Wallace *Parliamentary Counsel Office* [4.10 – 4.30]

Discussion [4.30 – 4.55]

Summary and close [Project authors]

[4.55 – 5pm]

Friday 27 September 2019

8.30 - 9am Registration, welcome tea and coffee on arrival

9am-10.30am

Discussion points from Day 1, orientation for Day 2	[9.00 – 9.05]
AI in the legal profession	
Jean Yang <i>McCarthy Finch</i>	[9.05 – 9.25]
Geoffrey Roberts / Matthew Bartlett <i>Citizen AI</i>	[9.25 - 9.45]
Mary Ollivier <i>Director, Regulatory, New Zealand Law Society</i>	[9.45 –
10.05]	
Discussion	[10.05 – 10.30]

10:30 – 11:00 Morning Tea

11.00am – 1pm

AI and Health / Elder Care

Pieta Brown <i>Orion Health</i>	[11.00 – 11.20]
Elizabeth Broadbent, <i>University of Auckland</i>	[11.20 – 11.40]
Angela Ballantyne, <i>University of Otago</i>	[11.40 – 12.00]
Discussion	[12.00 – 1 pm]

1 pm – 2 pm Lunch (provided)

2 pm – 3.30 pm

Digital assistants / digital counsellors

Jim Warren <i>University of Auckland</i>	[2.00- 2.20]
Gareth Cronin <i>AmbitAI</i>	[2.20 - 2.40]
Discussion	[2.40 – 3.15]

3.15 pm – 3.45pm Afternoon Tea

3.45 pm – 5 pm Digital assistants / digital counsellors continued

Lena Waizenegger <i>Auckland University of Technology</i>	[3.45 – 4.05]
Hazel Bradshaw <i>Department of Internal Affairs</i>	[4.05 – 4.25]
Discussion	[4.25 – 4.55]
Summary and close [Project authors]	[4.55 - 5.00]

Participants

Gordon Anderson is Professor of Law at Victoria University of Wellington and is a leading authority on labour and employment law. He is one of the authors of the leading commentary on employment law, *Mazengarb's Employment Law* (looseleaf, Lexis:Nexis) and the author of *Reconstructing Labour Law: Consensus or Divergence?* (VUP 2011), an account of the legal nature of labour reforms over the last four decades. His most recent book, *The Common Law of Employment*, written with Douglas Brodie and Joellen Riley, was published by Edward Elgar in 2017. Gordon was the lead editor and a contributor to *Transforming Workplace Relations in New Zealand 1976-2016* (VUP 2017). Gordon has written numerous academic articles on a variety of aspects of New Zealand employment law particularly on personal grievances, the legal and industrial relations restructuring of the last two decades and the introduction of a good faith obligation into New Zealand labour law in 2000. His current research interests focus on the changing character of labour regulation and the future of labour law. Gordon has represented various clients in employment related matters and he has provided policy and legal advice on legislation and labour law reform. He is currently the chair of the Ministerial Taskforce reviewing the Holidays Act.

Angela Ballantyne is an Associate Professor in the Department of Primary Health Care and General Practice (Wellington) and the Bioethics Centre (Dunedin). Angela's research interests include exploitation, research ethics, the ethics of pregnancy and reproductive technologies, and secondary use research with clinical data. She has worked in schools of Medicine, Primary Health Care and Philosophy in New Zealand, Australia, England and the United States; and as the Technical Officer for Genetics and Ethics at the World Health Organization in Geneva. In 2018 and 2008 she was a Visiting Scholar at the Yale University Interdisciplinary Center for Bioethics.

Matthew Bartlett is an Executive Director of Citizen AI and a Director of Loomio. He has a background in book publishing, web development and online legal content.

Dr. Matt Boyd holds a PhD in philosophy from Victoria University of Wellington and an MBChB from the University of Otago. He conducts bespoke health and technology research and has published on Artificial Intelligence and New Zealand policy in *Policy Quarterly*. Matt was a major contributor to the AI Forum NZ's research report 'Towards Our Intelligent Future' including authoring the Forum's soon to be released on AI and Health Care in New Zealand. Matt has previously conducted health technology assessment for the National Health Committee, and worked at the University of Auckland's Centre for Medical and Health Sciences Education. Matt has interests in public health research and research that helps to mitigate catastrophic risk.

Dr. Hazel Bradshaw is a game designer and academic who leads the Emerging Technologies work-stream within the Department of Internal Affairs, NZ Government. Hazel holds a doctorate degree in Human Interface Technology, specialising in design systems for serious-games, a Masters degree in Design, Strategy and Innovation and has held national and international academic positions, lecturing in creative technology, game-design systems and strategy and innovation processes. Hazel has published numerous academic papers and industry articles, outlining tangible design approaches for engaging broad groups in solving complex-problems, via the application of new technologies. Hazel focuses on building strategies for the tangible adoption, leveraging and navigation of the emerging technologies landscape; of which Artificial Intelligence is prominent, with the intent to prepare government for the impacts on society and public service delivery. These artefacts are drawn from a range of emerging technologies, such as Spatial Computing—for instance, Virtual and Augmented Reality, and Artificial intelligence tools, like Machine Learning (ML) and Natural Language processing (NLP). Hazel sits on the ‘Law, Ethics and Society’ working group of New Zealand’s AI Forum. Hazel and her team have contributed the ‘Legislation as Code’ use case assessing the ethical implications of legislation as code in an AI enabled democracy. Hazel is also conducting experimental work on how predictive algorithms are driving the erosion of social norms in an online setting. This work involves a Virtual Reality experience contextualising how predictive algorithms, (Narrow AI) optimised for ‘monetized clicks’, is driving the agenda of far-right online hate speech.

Dr Elizabeth Broadbent trained as an electrical and electronic engineer at Canterbury University to pursue her interest in robotics. She then worked at Transpower, Électricité de Tahiti, and Robotechnology. After becoming interested in the psychological aspects of robotics and in psychoneuroimmunology, she obtained her MSc and PhD in health psychology, supported by a Bright Futures Top Achiever Doctoral Award. She received an Early Career Award from the International Society of Behavioural Medicine and Early Career Research Excellence Award from the University of Auckland. She was a visiting academic at the school of psychology at Harvard University and in the Program in Science, Technology, and Society at Massachusetts Institute of Technology in Boston, USA. In 2017, she returned to Boston with a Fulbright award to study companion robots for four months. Her current research interests include how stress affects our health, how our body posture affects our mood, interventions to help patients make sense of and cope with illness, and human-robot interaction in health contexts. She is particularly interested in the emotional connections we form with robots, and how we can build emotional intelligence and empathy skills in robots.

Pieta Brown is a Senior Data Science Consultant at Orion Health where she works alongside the machine learning research and product teams to deliver tailored client solutions in healthcare. Pieta is driven to see data science deliver real value and improved outcomes and brings a unique background combining law, management consulting and data science to this challenge. Pieta has comprehensive knowledge of data-driven solutions across the

Healthcare, FMCG, Financial Services and Telecommunications sectors and is passionate about building useful and beautiful data products in New Zealand.

Gareth Cronin is Chief Technology Officer of Ambit, an AI-enabled conversational software platform, looking after product and technology vision, strategy, and execution. He is also an Executive General Manager at Xero, the cloud accounting and small business platform. Prior to founding Ambit and joining Xero, Gareth consulted for Air New Zealand and Vista Entertainment, led engineering at crime analytics software business Wynyard Group, at health software creator Orion Health, and at manufacturing software vendor Kiwiplan. Gareth serves on the board of the IT Professionals New Zealand's (ITP) national tertiary degree accreditation programme under the international Seoul Accord, the steering board of the NZQA mandatory review of IT qualifications, and the advisory panel for the Auckland ICT Grad School. In the past he has been a member of the ITP Auckland branch committee and the IT advisory panel for Manukau Institute of Technology. Gareth left an earlier career as a piano teacher and musician to graduate with a BSc(Hons) in Computer Science, work as a software developer, and later complete an MBA at the University of Auckland.

Colin Gavaghan is an Associate Professor in the Faculty of Law at the University of Otago. He is the first director of the New Zealand Law Foundation sponsored Centre for Law and Policy in Emerging Technologies. The Centre examines the legal, ethical and policy issues around new technologies. In addition to emerging technologies, Colin lectures and writes on medical and criminal law. He is deputy chair of the Advisory Committee on Assisted Reproductive Technology and a member of the Advisory Board of the International Neuroethics Network. He was an expert witness in the High Court case of *Seales v Attorney General*, and has advised members of parliament on draft legislation. He is co-Director of the Otago's Centre for AI and Public Policy and co-principal investigator in the AI and Law in New Zealand project.

Toby Gee is a barrister and mediator at Lambton Chambers in Wellington. He qualified in law after studying mathematics and philosophy at Cambridge University. From 1993 until 2013 he practised as a barrister (England and Wales) at Crown Office Chambers in London. He specialises in product liability (medical and non-medical), insurance and professional risks, and medico-legal issues, among other civil disputes. He has published articles and presented seminars on managing cyber risks. He was a panellist at New Zealand's National Cyber Security Summit in 2015 and for the New Zealand Insurance Law Association in 2018, and has contributed to the development of New Zealand's national Cyber Security Strategy. He has three children and enjoys real world activities such as ski mountaineering and singing.

Elizabeth George is Professor of Management in the Graduate School of Management of the University of Auckland. She has held academic positions at the Hong Kong University of

Science and Technology, Australian Graduate School of Management, University of Queensland and Western Michigan University as well as visiting positions in Duke University and the Indian School of Business. She has a Ph.D. in Organization Science from the University of Texas at Austin. She has an active research interest in nonstandard work arrangements and diversity in the workplace. Her work has been published in major international academic journals such as *Academy of Management Review*, *Academy of Management Journal*, *Administrative Science Quarterly*, *Journal of Applied Psychology*, *Organization Science* and *the Academy of Management Annals*. In addition, her research has been used by the *International Labor Organization* and the *US Society for Human Resource Management* to help inform public policy and management practice.

Avalon Kent is the Legal Officer at the New Zealand Council of Trade Unions. She has over 10 years' experience in industrial, employment, anti-discrimination, human rights, workplace injury compensation policy and law. She has a particular interest in pay equity and the regulation of non-standard & precarious employment. Avalon has a BA(Hons)(Industrial Relations), LLB & LLM (First Class).

Alistair Knott is an Associate Professor at the Department of Computer Science in the University of Otago, New Zealand. He studied Psychology and Philosophy at Oxford University, then took an MSc and PhD in AI at Edinburgh University. Ali has worked in AI for 25 years, focussing on models of natural language processing, human-computer dialogue and neural models of language and memory; he has published over 100 papers on these topics. He also works for the Auckland-based AI company Soul Machines, where he is implementing the embodied model of language developed in his book *Sensorimotor Cognition and Natural Language Syntax* (MIT Press, 2012). He is co-Director of the Otago AI and Society discussion group and co-principal investigator in the AI and Law in New Zealand project.

Joy Liddicoat is a lawyer whose primary research interest is human rights and technology. Joy is an Assistant Research Fellow on the AI and Law Project. Prior to joining the Project Joy was Assistant Commissioner at the Office of the Privacy Commissioner, managing policy and technology research and investigations into interferences with privacy. Between 2011-2014 Joy coordinated a global campaign *Internet Rights are Human Rights* advocating in the United Nations Human Rights Council, developing curricula and publishing related research. A Human Rights Commissioner for eight years, Joy was responsible for research on women's rights, national human rights institutions and led the Commission's 2010 inquiry into the experiences of transgender people in New Zealand. Joy is Vice President of InternetNZ <https://internetnz.nz/> which is responsible for domain name policy for the country code .nz.

James Maclaurin is a Professor in the Department of Philosophy and Associate Dean for Research in Humanities at the University of Otago. His MA in biological applications of mathematical information theory is from Victoria University of Wellington and his PhD in the philosophy of science is from the Australian National University. His research focuses on the relationship between science, public policy and ethics. His books include *What is Biodiversity?* (with Kim Sterelny, University of Chicago Press) and *A New Science of Religion* (with Greg Dawes, Springer Science). He has also published on philosophical methodology and on the application of evolutionary science in economics and computer science. He is co-Director of the Otago's Centre for AI and Public Policy and co-principal investigator in the AI and Law in New Zealand project.

Rakesh Mistry is a Product Manager at Straker Translations. Straker Translations with over 150 staff in 9 countries is one of the world's fastest growing translation companies. In his role Rakesh is part of the team at Straker Translations setting strategies, roadmaps and scoping out application developments to deliver high-quality translations over a technology-based platform. This includes discovering how Straker Translations can better utilise Machine Learning and AI to improve the overall delivery of our translation services. Straker Translations is a listed company on the ASX.

Paula O'Kane is a Senior Lecturer in human resource management. Her research spans many boundaries in the HR field but most recently she has been exploring the future of work, looking at the skills needed in the New Zealand context and how human resources and people practice might be impacted by automation and performance monitoring.

Mary Ollivier is the New Zealand Law Society's Director, Regulatory. Mary was admitted to the High Court in December 1991. She has practised in law firms in Auckland and Wellington and overseas. She was previously the NZLS acting Executive Director from January 2018 until April 2019. Mary is involved in all legal regulatory matters including at an international level and in relation to lawyers becoming reporting entities under the anti-money laundering legislation from 1 July 2018 and is a current member of the Council of Legal Education.

Jenna Riddle is a partner at Dunedin law firm Cook Galloway Allan specialising in alternative dispute resolution. She has a broad range of litigation experience across a variety of fields, including: general commercial litigation, insurance, employment, health and safety, transport, regulatory and criminal and construction.

Geoffrey Roberts is an Executive Director of Citizen AI and was previously General Manager of Community Law Wellington. He has a background in not for profit management and access to justice initiatives.

Paul Roth is a professor of law at the University of Otago, and specialises in employment law and privacy law. Since 1994, he has written the LexisNexis looseleaf commentary *Privacy Law and Practice*, and has written extensively on both employment law and privacy law over the years. He has also practiced in these areas of law, and has worked on various projects for the International Labour Organisation, the European Commission, USAID, New Zealand government departments, as well as other agencies.

Diane Ruwhiu is a Senior Lecturer in Management, with research interests in understanding the different modes and practices of Māori economy and enterprise. Her recent work engages with conceptions of indigeneity in work, organisation and management across fields of critical management, entrepreneurship, tourism and gender.

Dr Jeanne Snelling holds a joint position as Lecturer at the Faculty of Law and at the Bioethics Centre at the University of Otago. Jeanne first joined the Faculty of Law in 2005 as a Research Fellow on the Multidisciplinary NZ Law-Foundation-sponsored Human Genome Project, after which she completed a PhD focusing on issues arising from Human Reproduction. Jeanne currently lectures Law and Medicine; Bioethics and the Life Sciences; convenes the Masters of Bioethics and Health Law; and lectures on health law and ethics to undergraduate medical and allied health students at the Otago School of Medicine. Jeanne's research interests include health law, particularly the ethical and legal implications of genetics and biotechnology, and criminal law.

Grace Smart is a policy advisor working on Economic Development policy at the Ministry of Business, Innovation and Employment. Currently, Grace is working in partnership with the digital technology sector to create an Industry Transformation Plan for the sector. This work will include consideration of the opportunities and challenges presented by technologies such as Artificial Intelligence to the tech sector and the wider economy. Grace has previous experience working on policy related to the manufacturing and screen sectors, and in MBIE's Just Transitions Unit.

Lord Thomas of Cwmgiedd was the Lord Chief Justice of England and Wales between 2013 and 2017. He is an internationally respected contract scholar and theorist of private law, and has a particular interest in aspects of IT, digitalisation, data and AI – particularly their effect on law, legal practice and the courts. Lord Thomas is the New Zealand Law Foundation's 2019 Distinguished Visiting Fellow. He practised at the Commercial Bar in London from 1972 to 1996, becoming a QC in 1984. He was appointed to the High Court of England and Wales in 1996. Lord Thomas was successively a Presiding Judge in Wales, Judge in Charge of the Commercial Court, the Senior Presiding Judge for England and Wales, a Lord Justice of Appeal, and President of the Queen's Bench Division.

Dr Lena Waizenegger is a lecturer in Business Information Systems at the Auckland University of Technology, New Zealand. Lena received her PhD in Information Systems from the University of Innsbruck, Austria in July 2017. Her main research areas are (a) ubiquitous connectivity and team collaboration (b) conversational agents and (c) digital disconnection. In her research on conversational agents, she investigates how the technology is perceived by human employees that work “alongside” their new digital colleagues and the effects of conversational agents on the customer experience. She further explores how robots as team leaders that use different control styles are perceived by human team members and affect their work autonomy, job satisfaction and expected task performance. Her research has been published in various peer-reviewed Journals such as the International Journal of Knowledge Management (IJKM), the Journal of Travel Research, Vocations and Learning, and Cutter Business Technology Journal as well as conference proceedings such as ICIS, HICSS, and ECIS.

Richard Wallace is Parliamentary Counsel and Drafting Team Manager, Resources and Treaty, at the Office of Parliamentary Counsel. Richard has worked at the Parliamentary Counsel Office as a legislative drafter since 2003 and was appointed as a Drafting Team Manager in 2012. He established and led the Access to Secondary Legislation Project, which aims to collect and publish New Zealand’s tertiary legislation on the NZ Legislation website, until September 2017. Richard has been involved in researching, developing and testing ways in which technology could change the manner in which legislation is drafted, made available, accessed and used. His particular areas of focus have been on producing rules as code, and developing the Better Rules methodology. Richard has a background in commercial law, having practised in New Zealand and England.

Sara Walton is an Associate Professor who teaches and researches in the area of sustainability and business often with a future focus. She is the director of the Master of Sustainable Business and chairs the Postgraduate Committee and postgraduate processes in the department. Sara has contributed towards knowledge in environmental and social entrepreneurship, innovation for environmental sustainability, sustainable transitions, environmental conflicts and understanding work in changing futures. She is part of the Otago Futures Research team that systemically explores narratives of the future of work, producing reports on the future of work in Dunedin (2013), for the ICT industry in Dunedin (2018) and New Zealand High Value Manufacturing (2019).

Jim Warren is a Professor of Health Informatics at the University of Auckland, based in the School of Computer Science. He specialises in design and evaluation of information systems to support long-term condition management. He has worked extensively with the National Institute for Health Innovation (NIHI) at the University’s School of Population Health. With NIHI he has consulted on implementation of New Zealand’s national health IT plan in areas including electronic referrals, shared care planning systems and information architecture.

His current work includes statistical modelling and machine learning to improve understanding of cardiovascular disease risk in New Zealand and development of an IT platform for coordinated screening and e-therapy to support mental health of New Zealand youth. This latter work is sponsored by the National Science Challenge, 'A Better Start – *e Tipu e Rea*' and CureKids. For over 10 years he has worked with the Department of General Practice and Primary Health Care on a consumer-operated health and lifestyle e-screening tool, eCHAT (the electronic Case-finding and Help Assessment Tool) and in recent years this tool has been expanded and trialed for young people (as YouthCHAT). He has been engaged with HINZ as three-time conference chair (2006-2008) and overall chair of the organisation (2008-2010) and has served on New Zealand's health informatics standards body. He is a Foundation Fellow of the Australasian College of Health Informatics. His degrees are in Computer Science and Information Systems from University of Maryland. After completing his PhD he worked for one year at The American University in DC and from 1993-2005 for University of South Australia in Adelaide.

Jean Yang is Chief Operations Officer and Vice President Legal Services of McCarthy Finch. Jean is helping reshape how law is practiced at legal AI business, McCarthy Finch. Originally a solicitor at Minter Ellison Rudd Watts, she was the first lawyer to join McCarthy Finch and has since built a team of Legal Engineers who work on the intersection of law and technology. Jean is a founding executive of LegalTechNZ, whose mission is to facilitate the development and adoption of legal technology in New Zealand.