INSTITUTE OF
PUBLIC ADMINISTRATION
AUSTRALIA
ACT

02 6154 9800
PO Box 4349 Kingston ACT 2604
admin@act.ipaa.org.au
www.act.ipaa.org.au

ABN 24 656 727 375

TRANSCRIPT OF EVENT

EMERGING TECHNOLOGY: AI IN AUSTRALIA AND THE FUTURE OF DECISION MAKING

PROFESSOR MICHAEL MILFORD (PANELLIST)
Deputy Director
QUT Centre for Robotics

DR CHENG SOON ONG (PANELLIST)
Principal Research Scientist
CSIRO

CLAIRE CLARKEE (PANELLIST)
Director of Methodology Futures
Australian Bureau of Statistics

DR PAUL WONG (PANELLIST)
Senior Fellow, 3A Institute
The Australian National University. The panel discussion was Chaired by

BELINDA DENNETT (HOST)
Corporate Affairs Director
Microsoft Australia

17 November 2020

Held in partnership with the Graduate Data Network.

Enquiries should be directed to Caroline Walsh on 0413 139 427 or at caroline.walsh@act.ipaa.org.au

BELINDA DENNETT: Good afternoon, everyone. And welcome to today's event. The *Emerging Technology: AI in Australia and future of decision making* session. My name is Belinda Dennett and I'm the Corporate Affairs Director at Microsoft Australia, and I'll be your chair for today's event. And I'm pleased to open this digital event coming to you live from the IPAA offices in Barton. We have some panellists, including myself joining remotely, making this a truly digital event. Today's event is delivered in partnership with the Graduate Data Network.

I'd like to begin today by acknowledging the Ngunnawal people, the traditional custodians of the land on which this video has been broadcast from. And the Wurundjeri people, the traditional owners of the land where I'm coming to you from. We acknowledge and respect their continue in culture and the contribution they make to the life of this city and this region. And I'd like to acknowledge and welcome any Aboriginal and Torres Strait Islander people who may be tuning into today's event and to the elders of all lands this broadcast reaches.

We have a great panel of speakers today, and I'll introduce them to you shortly. To our audience, I'd like to invite you to submit a question virtually, which you can do at any time throughout our panel discussion. Please provide your name and agency when providing the question and we'll endeavour to answer as many of those as we can during the Q&A. The overview of today's event is we're going to have the panel discussion followed by audience Q&A. Today's event is being recorded and will be made available to view online as part of IPAA's post event resources page.

Now it's my great pleasure to introduce our panellists. We have Professor Michael Milford, the Deputy Director of the QUT Centre for Robotics, a Professor at the Queensland University of Technology, Microsoft Research Faculty Fellow, and the Chief Investigator at the Australian Centre for Robotic Vision. Michael's research models, the neural mechanisms in the brain, underlying tasks like navigation and perception to develop new technologies such as all weather, anytime positioning for autonomous vehicles. He is passionate about engaging and educating society about the opportunities and impacts from technology such as robotics, autonomous vehicles and artificial intelligence.

We have Dr Cheng Soon Ong. Cheng is a Principal Research Scientist at the Machine Learning Research Group at Data 61, as well as the Director of Machine Learning and Artificial Intelligence, Future Science Platform at CSIRO. Cheng is interested in enabling scientific discoveries by extending statistical machine learning methods, and is also an Adjunct Professor at the Australian National University.

We have Claire Clarke, Claire's a Principal Data Scientist and the Director of Methodology Futures at the Australian Bureau of Statistics, and the ABS Methods Architect. Her work includes the application of machine learning on data validation and missing data problems. The interpretability of machine learning and supporting innovations in organisations.

And we have Dr Paul Wong is a Senior Fellow at the 3A Institute at ANU, and is a data practitioner interested in the strategic use of data as an enterprise national asset. Paul's managed enterprise class systems, data, and has played a key role in the establishment of a national consortium to support digital research infrastructure. He's also contributed to the improvement of research data management practices at a national level.

So welcome to you all. I think the best way to get to know you a little bit is I'm going to invite each of you, each of our panellists to share some brief opening reflections. Michael, I might ask you to start us off.

MICHAEL MILFORD: Sure. Thanks Belinda, and well done to the organisers for putting together what I hope will be a really interesting panel discussion. I guess I should say upfront that overall I'm a cautious optimist about artificial intelligence and the effect it will have on society at large. I think it's the usual things of any new potentially transformative technology. It's balancing being careful, being cautious, being as certain as we can be about the unknowns with making sure we do explore the opportunities that AI affords us or could afford us.

Given this is a graduate audience, I guess you're looking forward to the rest of your careers and you could have quite long careers, many, many decades from this point onwards. And I know for people, even people who have been working in the field for several decades, trying to get on top of AI can be very overwhelming and daunting. And I guess my key takeaway there would be, don't be overwhelmed because everyone else is equally or more so overwhelmed. And you can choose sort of where your expertise is going to develop. 20 years ago you could argue the case that some people were world-class technical experts and simultaneously world-class experts in terms of regulation, ethics, and practice.

Nowadays just because modern artificial intelligence techniques are so incredibly sophisticated, it's hard to be both. So you get to choose where along that spectrum of technical expertise to sort of, the sort of house side of things you want to play and where you want to have most of your expertise. You're going to be learning about AI most likely for the rest of your careers. And I find that concept exciting and you have an advantage over the old people like myself, who have had to go back and refresh our understanding of modern artificial intelligence techniques with our old brains. You're taking it on with fresh brains and you're very well equipped to sort of keep track of this extremely rapidly moving field.

The other thing I'd mentioned too, that we see come up a lot in this sort of grad type seminars is it's been tempting for experts and commentators to try and simplify AI down to sort of core concepts, core rules of thumb. And they're very useful up to a certain point. But I'd urge people, especially people who are going to be active in this field for potentially four or five decades into the future to treat very cautiously some of these sort of absolute statements you'll see. These are very comforting, these statements, like AI is great at highly automated repeatable processes, but it will never be emotionally perceptive. It will never have empathy.

I think when people say that, what they're really saying is that with current techniques, those very human like problems are extremely challenging and we don't know whether we can solve them yet, but on a 40 or 50 year timeline, I think it's important that we keep an open mind that we may make significant progress on some of these things that we consider currently uniquely human like. It's not something to be scared of or to be afraid of, but it's going to influence I think how we use and leverage AI over the rest of our lives. And thank you for having me on the panel.

BELINDA DENNETT: Thanks Michael. Cheng, we might throw to you.

CHENG SOON ONG: Thank you. Thank you to organisers for inviting me. And thank you to you all online for spending your lunchtime listening to me. At heart, I'm very much a scientist and I like to think of how AI can augment what scientists do. And if you think back about the scientific method, it's roughly 400 years old today invented by Francis Bacon in 1620 more or less. The process of science is this combination of going from data to knowledge, and using the knowledge you have to find new data. This process of going from data to knowledge can be thought of very much like an observation process.

And you can think of AI as contributing to the observation process. For example, AI could be used to look at satellite pictures and say something about the landscape of Australia. Now, what I find a lot more interesting. And I think this falls to the theme of today's conversation is this idea of, if I have some knowledge about the world, what type of data should I be collecting? Now, this has multiple dimensions to it. You can think of it purely from a technical point of view, to maximise the kind of information you want to collect. You might want to think of it from an ethical point of view. Some types of data may not be good to collect. So you can think of it in multiple forms and that's the loop, the closed loop of scientific work.

Now, one thing I've observed a lot of in my work is the fact that as the AI becomes a more ubiquitous technology, that you often need to be an expert in some other area to make an impact with AI. Concretely, if you happen to be a biologist or something, and you're working in an agricultural system, then it's not good enough to be the world expert in AI, you actually need to know about the system you care about. And this, I think has this key point, which Michael already mentioned, that you're going to have a whole career in front of you and you want to be an expert in some domain. There's a phrase that gets used a lot, which I don't particularly like, but it's a good analogy that AI is the new electricity.

We don't talk about people who are experts in electricity. I mean, they are people who are electricians, but by and large most of us use electricity to do something else. And I can see as the community develops, we're going to not talk about AI enabled something, it's going to be just agriculture. Thank you.

BELINDA DENNETT: Thanks, Cheng. Claire, we might ask you to share your opening thoughts.

CLAIRE CLARKE: Like the others, I'd like to thank everyone for inviting me here and for this opportunity to talk to you about, there's lots of very exciting developments going on in this space. I thought I might talk a little bit about some of the ways that both the ABS and other statistical offices are looking at using AI at the moment, because I think the very broad sort of themes that are probably very resonant with not just other statistical agencies, but just other agencies generally. There's kind of three sort of broad ways that we're looking at using sort of machine learning and AI at the moment.

One is the volumes of data available in the world at the moment are increasing just by orders of magnitude all the time. Long gone are the days when your only source of information about Australia was sending the census collectors out on horseback to ... And we're definitely getting to the point where it's actually really, really hard to continue just with our standard sort of manual processes of trying to turn these volumes of data into useful information. And there's sort of traditional statistical techniques that we've used in the past. They kind of stumble a bit in the face of very frequent data, high volumes of data. If we want to make use of the opportunities presented by these new data, and if we want to make the best use of them, then increasingly we need to turn to machine learning and artificial intelligence.

There's also, I guess, is we would look at ways that we can make our business more efficient. As I said, there's still a lot of quite tedious and time consuming work I guess, that goes on behind the background sort of in our business and probably in lots of other people's businesses as well. And the opportunity to free up our skill staff from doing that kind of boring stuff and getting to do the more exciting stuff. It's a really good use of artificial intelligence I think. Just the other thing is really that these are a really useful set of tools that we can add to our standard set of statistical tools. And even if we're not looking at interesting new data sets, satellite data, social media data, all that kind of thing. The opportunity to sort of draw new insights from the old stuff that we have, because we've got this new set of techniques.

BELINDA DENNETT: Thanks Claire. And Paul, I'd love to hear your opening thoughts.

PAUL WONG: First, I would like to acknowledge the Ngunnawal people and play my respect to elders past, present and future. I'd like to open with some provocative remarks, three, possibly four if time permits. First, I want to comment on emerging technologies. The first question I would like to ask is emerging from what and where. That question is important because technology has a history. And the kind of history that we're looking at is the history of AI. If you know that history, AI has been around for more than 80 years, is not a recent invention. It came about because of the work of mathematicians, Alan Turing, Alonzo Church, their work on computations, literally revolutionise the way we think about our computers. And without computers there can be no AI as such as we know it. History is important in understanding technologies.

The second provocation is AI in Australia. I want to suggest that we look more broadly to beyond Australia, because AI is an international endeavour. Australia is doing really well in terms of it's R&D. If you have a look at the Australian Research Council, excellent research Australia assessment result. The field of research, artificial intelligence and image processing are rated very highly. Of the 30 institutions, universities rated, seven institution rated as well above world standards, at 11 rated above world standards. We're doing really well in R&D and in artificial intelligence. But the question is, are we harnessing those opportunities and turning those research findings into tangible product or services that benefit the community in Australia. There's a bit, I think a bit of gap in translations of research.

And the third provocation is about the future of decision making. Again, question, what kind of decision are we talking about? Who will be affected by these decisions? What kind of effect will these decisions produce and are they safe decisions? Pack into that title emerging technologies, AI in Australia, and the futures of decision making. I think is a lot of questions that we need to think deeply. And my prediction is that these questions will not go away in the next five to 10 years. We will continue to have to come back to these questions, ask ourself, what are we doing? Why are we doing this? And is it for the benefit of Australia.

BELINDA DENNETT: Right. Thank you, Paul. Really fascinating insights to open there with. And I guess from my perspective as a Victorian who has just come out of a lockdown for 120 days. And just the discussion, the mainstream discussion around data science algorithms, modelling around the decisions that were being made around COVID and lockdowns and restrictions. Like everything, we've been in my industry talking a lot about the role COVID played in accelerating take adoption and take update, and sorry, adoption and uptake. And I feel like COVID's probably had a role on the way AI has entered into the mainstream awareness. Let's go to some questions to open the panel discussion. And perhaps I'll start with Michael, and if you could share with us your thoughts on the influence AI might have on the future of public service delivery.

MICHAEL MILFORD: Sure. Maybe I can speak to some general principles because there are members of the panel who have more specific examples. But if we look at how AI and related fields and machine learning and all sort of peripheral fields are translating the most into actual practice, it's typically in areas where you are making large numbers of decisions or actions rather than one, because if your system that you deploy makes the wrong decision, the outcomes are not catastrophic. It just comes out as a sort of number averaged across a large number of decisions. If you look at other fields like autonomous vehicles or flying planes, or AI is sort of deployed to various extents, the big hold-ups in those other fields are you don't want a system that you can't prove completely in charge of decisions that could have immediate catastrophic effects.

The other problem with introducing AI, especially modern AI, I define modern AI as sort of the last five or 10 years of development of modern artificial intelligence techniques, which are typically far more powerful, but also a lot more opaque in terms of understanding what they're doing. They have the most potential upside in terms of improving what we already do manually, or through traditional sort of statistical techniques, but they're also the hardest to deploy because we don't understand exactly what they're doing, and we don't understand the time lag in the effects of what they do. Some of these effects good or bad may not become apparent until five or 10 years down the track. And by then you've already screwed it up.

I think in terms of the short and medium term areas where it'll be deployed, it will be in systems where you can audit the outcome where you can characterize quantitatively, whether your system is really doing what you think it's doing. Those are the sort of areas where I think the most safe gains can be made in the short and the medium term, while the researchers work on ways to better understand and characterise these very complex sort of black box systems, so we can make those sort of transformative gains in the longer term.

BELINDA DENNETT: Great. Thanks, Michael. Perhaps if I go to Claire, and Claire, do you want to talk a little bit about how AI might influence the future of decision making processes amongst the ABS leaders?

MICHAEL MILFORD: You're muted Claire.

CLAIRE CLARKE: Of course I am. I'm sorry, I should also point out that with exquisite timing, the council has decided to move the verge outside of my house at the moment, so I apologise if a little bit of noise comes through.

BELINDA DENNETT: There's always a lawn mower, I love it, it's very Australian.

CLAIRE CLARKE: I think perhaps to sort of follow in a little bit from what Michael said. I think I see potential very much for a sort of supported decision making. The more complex our systems get and the more data and information we have available. I think the role of AI will be to make that manageable and consumable for a human so that they can look at that collated set of information from a range of different sources. And whether that comes with some kind of probability of this happening or possibility of this. And you kind of, it sort of supports that decision making that you might then have to do from that.

BELINDA DENNETT: Thanks Claire. Look, I think that's interesting. Again, if I use my COVID example, I think that's perhaps where some of the failure was that there was a lot of talk about the data and the science and the decision making. There perhaps wasn't a lot of explaining what that was telling people and how that was informing decisions. And I think that does become ... If people have to trust in AI, then they need that bit of explanation as well. Perhaps if I go to Paul, given at the university and the 3A does a lot of work with students, interested in what skills, tools and resources do you think graduates should be equipped with in order to fully take advantage of the potential of AI?

PAUL WONG: My sense is that systems are getting ever more complex. And in order to have a grip on these systems, I think we need to have a more holistic way of understanding systems, how different components may be connected, how people interact with these components. As many nodes when system are complex, they can produce unintended consequence. And that's I think part of the reason why the 3A Institutes are concerned about the safety of mass deployment of these systems, AI enabled systems. And that's why we're being deliberately provocative in asking those questions, who'll be affected by these decisions when these semi-autonomous systems are making decisions on behalf of us. I think to answer your questions, the ability to think more broadly, to think about how systems are connected and how we are place in part of the ecosystems of these systems is absolutely critical.

BELINDA DENNETT: All right, thank you. And Cheng, sort of following on from that, given our graduates are probably as Michael alluded to, going to be at a great advantage in that they are going to be learning AI from the start of their careers. How do we bridge that divide between what these graduates may understand and the senior leaders who this maybe new and really turning decision making processes on their heads from how they've always done their roles?

CHENG SOON ONG: I think it's quite important like always, and it hasn't really changed is communication. I think I'm particularly aware of this working across disciplines that if you have a new technology, it is always very attractive to just deploy it as widely as you can. But in some sense, and in some sense we need to move with the times such that we adopt these new technologies to make the best use of them. Now to do that I think it's particularly important that the young people who understand the technology try and communicate the benefits and the impacts of this technology, and to communicate that often in some other domains.

In the public service, we do a lot of efforts with respect to deploying systems and write policies that intervene not to improve the lives of Australians. And I think it's important to understand, if we put this new technology in, you want to communicate to your senior leaders that, what are the impacts of this technology? Who is it going to be affected? What are the quality controls that we're going to do? How are we going to change our regulations to address the fact that we have these new technologies? Those questions are always there. And I do agree with all the other panellists here, these questions are always there. If you do understand AI, it's up to you to communicate what you know in the language of the other person.

BELINDA DENNETT: Great. Thank you. I might just throw back to Michael. This is one of my favourite topics when we talk about AI, is how do we build trust in the technology. And given your area of robotics and automation and automated vehicles, I'm really keen to get your views on how important it is to build trust and how you think we go about doing that.

MICHAEL MILFORD: Building trust is supremely important because although AI is technically making great leaps and bounds and has over the last decade, I think it's becoming abundantly clear that a large percentage of realistic usage scenarios are really going to involve human machine teaming to some extent. It's incredibly hard to entirely automate a system in a way that works and is safe and is trustworthy. It's far easier to do 95% of it and have a human looped in at the appropriate times. Trusted autonomy, which the centres throughout Australia that are focusing on trusted autonomy is going to be vital because you're going to want the humans who are interacting with these systems who are maybe commanding them, or maybe using information from them to make decisions. They're going to need to trust these systems. And part of that trust is going to be around topics of explainability. And this is really, really tough, right.

Because let's take health care for example. There's been a lot written about the fact that modern AI systems combined with sort of big data, they're going to be able to put together patterns and trends that no human or even human team could ever do on a global scale. Now, the machine is going to be able to explain that at a high level, by saying, "Look, I found some very significant seventh order relationships in the data, and that's why I'm recommending this course of action." But in terms of actually explaining the details of that, the reason humans haven't done that before is that sort of information is very inaccessible to a human mind. There's still a lot of work that needs to be done on trust, on explainability. But I think the progress is good. And as I said before, I think we can work on those problems in the background while we make sort of incremental advances in the sort of more short term implementations of AI that can help us.

BELINDA DENNETT: Great. Thank you. We have a few audience questions coming in. I think this is a good one and perhaps we skipped over this at the start. I might throw to Claire on this one, Rachel from Services Australia asks, "How do we define, what's your definition of AI and how much does it overlap with process automation?"

CLAIRE CLARKE: That's an easy question. It is actually quite a difficult place to draw the boundary in some ways, if you go along to machine learning course, that sometimes some of the first things they'll start talking about is things that a statistician report would think of as, that's just ordinary statistical theory. I think that it's probably just maybe a degree of complexity, but perhaps more so like taking a few steps away from human involvement as well. If you think about process, so with that comparison with process automation, you can automate some things very simply, you can just write a few lines of computer code to say, "If this happens, do this, if this happens, do that." That's really just ... there's kind of one level of taking the human out of it. You've put the human who's come up with what those instructions are and put them in place.

But it's very sort of simplistic. I think a lot of, when people think about machine learning and AI, they think about taking several more steps away. It's not just the human saying, this is the process and these are the rules and this is the steps that you have to follow. But the computer itself is starting to decide what those things should be as well.

BELINDA DENNETT: All right. Thank you. Look, I guess, following on ... perhaps Paul this one to you. There's been discussion in Australia around whether we need a specific regulation for artificial intelligence, whether we need a specific regulator, whether we just need to enforce the laws we've got. How are you viewing sort of current public policy and the regulatory sort of framework around AI just in your views?

PAUL WONG: On the near horizon the particular kind of AI that is being deployed in the marketplace is very data-driven. That means that access to data is absolutely critical to deploy these services, these products. And I don't see a way out of regulations when we're talking about data, because we're talking about potentially data pertaining to personal information and so on collected through webcams or collect through CCTV. The other day when I visit my specialists, I had to sign a consent form because they had a CCTV in the office. So is pervasive precisely because data is the kind of the lifeblood of the current crop of AI.

Let me also make another comment that when we use the term AI, we almost as if we're using synonymous with machine learning. The actual fact is that AI is a very broad church, include many different analytical, many methodologies. Machine learning is one recent success of artificial intelligence. If other progress is made in AI, that doesn't require data to drive AI. Possibly regulations would not be as severe because of that. But just going back to the previous discussions about the role of decision making, once again, when the decision is made by machines, we need to have a way to assess the risks involved.

BELINDA DENNETT: Great. Thank you. I've got another audience question here, perhaps Cheng this one to you. Matt from the Victorian Department of Health and Human Services asks whether there are exemplars you're aware of where organisations have balanced the ethics of AI in practice?

CHENG SOON ONG: That's an excellent question, Matt. Firstly, I think there's what I'd like to call ethics washing going on. I think a lot of organisations, a lot of companies, they bring in ethics as something that's supposedly going to improve their processes. And there are many examples of companies who are trying to do things ethically, and I think the biggest partnership right now is called the Partnership for AI, which is a group of large companies in the world are signing up to perform ethically with respect to some code of conduct.

Now, it's important to remember though, that on one side there is this public relations exercise, which maybe I'm too pessimistic about this. But on the other side, there's what happens on the ground. And it's always unclear exactly what the impacts of these high-level statements are going to have on the ground. So if you're asking directly, is there this program that this company has done or this organisation has done that has made a difference, I think it's still too early to tell.

I think we adjust at the beginning of implementing widespread use of these decision making technologies, these automated decision making technologies in our society. And we don't know what the impacts are yet. And I'm just going to quote Michael Jordan, not the basketball player, but a statistician in Berkeley, a famous statistician, the AI revolution hasn't happened yet, right, it is just a beginning. We're just seeing the beginnings of the impacts of these. And so the impacts of ethics statements and things, I think the field is wide open. I encourage anybody to get in.

BELINDA DENNETT: Thanks Cheng. Claire, maybe this one's for you. Helen from the Digital Transformation Agency is interested in the quality of data and how this affects assumptions baked into AI and decision making, and how might we mitigate those unintended consequences?

CLAIRE CLARKE: This is a very important question. And it's certainly one that we're thinking about quite a lot. The quality of data that goes into decision making is obviously vital. So you really, if you don't want to be, I'm assuming, say for example, that you're making decisions involving the whole population, but you've missed some crucial sort of subpopulation, for example. It's really quite tricky, and it's kind of an active area of research I think because certainly in sort of traditional statistical approaches, you can handle some of those things. If you do actually end up with sort of some data that's got some missing values or something like that, there are ways that you can handle that. And the information that you get out of the end of things like standard errors and biases and things like that, can give you some information about how you're missing that.

And a lot of AI approaches don't necessarily incorporate those things yet I think, I would say that that's a yes. It's definitely something that we're thinking about. And lots of other people are thinking about. But it's a very important thing. And I think for the time being, we don't have the technical solutions to that necessarily yet. It is something that people do have to be very much aware of. A lot of methods I think, the garbage in garbage out kind of philosophy applies. If your daughter is nonsensical, you'll get a nonsensical answer. But there's potential to do a lot of damage with bad daughter in this space, I think. And until we have some ways of really clearly indicating something that comes out of the end of a decision making processes, that flows some bright red lights and says, "Hey, you haven't thought about this." We just have to be really vigilant, I think and really careful and sort of ask questions. Okay, so yeah, you've come to me with this thing you've done, you've done a model or whatever, you've got this prediction, you've got this recommendation. What was your starting point for that? What were the questions that you were asking at the very beginning and what was the data that you were using?

BELINDA DENNETT: Great. Thanks Claire. Question from Justin at the Australian Electoral Commission, and I might take the liberty of answering this one myself. Justin, you've asked how you would start a career in AI machine learning, just a little plug for a government industry initiative that we've done since COVID is that all the tech industries have worked with the department of industry to put together a skillfinder.com. And there's a whole lot of free courses and resources, many of them in data, data analytics, machine learning. That's worth having a look at. But you also ask about whether a PhD is what you would advise, so I'm going to throw to our university representatives. Michael, do you want to have a go with that one?

MICHAEL MILFORD: Sure, I'll get it started. I think the key to remember here is the answer to this question is very different to the answer even two or three years ago, and it's likely to be different in another two or three years from now. I guess to break it down into a few different components, firstly is you need to acquire the competency, the technical familiarity in the first place. And you can do that through university courses, although they tend to lag quite significantly because the field is developing so quickly. In AI, particularly the quality of online sort of formal courses is very good if you know which ones to go to. And there are literally millions of people self-educating themselves around the world for free.

Certification in our current education system is still important. You can definitely make your way into very exciting roles without a PhD. But realistically, at this current time, all other things being equal, certifications like a PhD can help get your foot in the door more quickly than without it. There was one more point I was going to make about, oh yes, seniority. The other thing you should look at is, although the start-up scene in AI in Australia is very fledgling. The numbers we have are very small compared to some other leading countries. Increasingly these start-ups just because they need top talent are becoming very creative in how they acquire talent. They're taking people from physics degrees or electrical engineering, undergraduate degrees, and training them up rapidly in-house over a 12 month period to become AI experts. Their traditional sort of rigid pathways are very much loosening somewhat, but they're not completely gone yet.

BELINDA DENNETT: Thanks, Michael. And Paul I might ask you if you've got anything to add to that from ANU 3AI perspective.

PAUL WONG: I always encourage people to do a PhD if they want to, if they have the appetite to do that. A knock about technical competencies, I think the space is broad enough that we can have a lot of different kind of expertise playing this space. Belinda, you mentioned earlier about regulations. There's space for people to think deeply about how do we set up a regulatory framework for AI. There are people who need to define policies. My sense is that the field is wide open and we will need a lot of people who have broad set of skill sets to contribute to the landscape.

BELINDA DENNETT: Great. Thank you. Another audience question, perhaps for Cheng, Leah at the Department of Health, is there any learnings that you can talk about around machine learning in healthcare data?

Dr CHENG SOON ONG: Well, that's a lot of learnings, but let me just pick on one that I'm quite passionate about. I think with respect to applying machine learning to healthcare, it's very important to remember that actually we don't have enough data and we will never have enough data. And the reason I make that statement is because we have this huge variety of conditions, people, backgrounds, environments that people live in. And in some sense every individual patient you see in a health system is a unique patient that you will never see again and you have never seen before.

And one of the key challenges in applying something like machine learning to health care is the fact that we don't have enough data about these very, very specific cases to make statements about these specific cases. And it goes to this deeper problem that very often the quality of our data is dependent ... Sorry, backwards. The quality of the machine learning method that we apply is very dependent on the kind of data that we have collected in the past. A very good example is the fact that most cars, so most safety systems of cars, I know it's not very related to health, they are designed for men who are 180 centimetres tall, because that's the size of the crash test dummy. And this is the kind of data we collect when we crash cars.

Now, if you happen to be a small Asian woman, then it is pretty tough for you being in that system. Now, the same thing applies, if you look at healthcare data, I think a lot of the data, a lot of the studies on health are done in, I'm going to call them caucasian countries, right? And so a lot of the outcomes we get in health are dependent on these. If you happen to be in some minority, and most of us are minorities in some way, then that general purpose data is not applicable to you in some sense. In that sense there's a huge challenge. And I see this also as an opportunity, because I think we have to think very hard about the fact that we never have enough data, and how do we deal with this if we want to do machine learning.

BELINDA DENNETT: Thank you. Lots of interest in the ethical questions. And also I think certainly one of the topics that I've noticed seems to be more prevalent in Australia than in many places. Perhaps, Claire, I'll throw this one to you. The workforce implications of automation in the public sector. I think robots are going to take my jobs. Do you have some thoughts to share on that kind of concern and fear amongst people?

CLAIRE CLARKE: Yeah. I mean, obviously it's a very real concern that people have. My feeling is that the robots are going to completely take my job is perhaps quite a long way away further than we think. But I think it's probably going to be more sort of along the lines of thinking about what parts of my job are going to be taken by a robot? What parts will be left for me or for other people I guess? And what does that mean in terms of the way that my job is going to evolve into the future? Will there be things I'll be spending more time on, things I'll be spending less time on? And how do you ... I think that there's a very interesting set of questions about how humans and computers sort of work alongside each other in ways that we can be most useful to one another in the future I think so. How does the computer help me and how do I give information to the computer to help it do its job as well.

I think in terms of alleviating those fears, I think, I mean it's always kind of difficult. But I think if you sort of talk about them in those terms a little bit, it might not necessarily be true for everybody to say that, well, the computer is going to take the most boring part of your job and leave you with the most exciting parts. There's probably a certain amount of privilege in thinking that. But I think the more kind of information that you can give people probably the better. And I think the more information that people can have about the kinds of things that computers are good at versus the kinds of things that humans are good at, even if that's an evolving kind of concept over time.

BELINDA DENNETT: Thank you. Thea from the Centre for Public Impact, this is an interesting one. I might go back to you Cheng. I think, again, I keep coming back to my COVID example, but when you think of the amount of data and the amount of research that will be done on this pandemic, when you look back to 1918 and the Spanish flu and the amount of data we had from that, I think it's quite stark. So Thea asks, "How can we use historical data effectively to make predictions in complex systems?" Cheng, do you have some thoughts on that one?

CHENG SOON ONG: I think in some sense we only have historical data, right? I mean, we never have data about the future, right? And in some sense, for dealing with this pandemic, we are actually often interested in something that's quite difficult to answer with current technology. We're actually interested in how we should intervene, what types of policies we want to do, right? And we build those policies and interventions based on data from the past. And this area of research is often called causal analysis. We were interested to do causal analysis.

Now, the tricky part about doing causal analysis from historical data is often we don't have control of how that data was measured. And this error is called causal analysis from observational data. Now, the gold standard in causal analysis is of course to do randomised control trials, but it's something that we cannot really do in an example like this pandemic here. But on the other side you've probably seen on the news that this is very recent press release, that Pfizer said, "Oh, we did these trials, and our vaccine looks like it's 90% efficient." Now, how do they do that? Right. I think working out how we can use historical data in a way that would allow us to make statements that are in some sense causal, I mean, this is a very deep philosophical question. What does it mean to cause something else, because there are feedback loops and things like that.

I'm not even going to try and answer that, but I think the challenge of using historical data to say something about the future actually requires often domain knowledge, whether it's domain knowledge and health or domain knowledge in public policy. Because this kind of focuses us down the path of causality, which is, statistics doesn't really have very good answers quite yet.

BELINDA DENNETT: Thank you. Another ethics and trust question, perhaps Michael, this one to you. Nick from the Department of Education, Skills and Employment asks, "Is there any incentive for large users of data driven AI to behave ethically in the long-term if there's no regulation in the first place?"

MICHAEL MILFORD: That's a good question. Look, I think if you want to encourage that sort of behaviour, it's going to have to be a multi-pronged approach. Regulation is going to have to be in there somewhere. People are going to push back against regulation because over-regulated systems will stymie progress. But you don't want unhindered progress with no thought of the consequences. One of the other touch points you have to influence large companies is their personnel. There've been a number of examples recently where large percentages of people at top tech companies have made a very public stand against some of their employees policies or what they were doing. It hasn't always been effective, but if you can make sure that the new talent, so all the grads coming through the system have a sort of deeply informed understanding of the potential ethical and societal impacts and considerations of the technology that they may be working on, or part of, that's going to be a very good lever to try and influence the behaviour of these behemoth, large corporations, because you will only get so far with sort of regulation and carrot and stick type exercises from government.

BELINDA DENNETT: Thanks, Michael. And probably our last question is, let's go to Paul, Nicholas from the Australian Competition and Consumer Commission, "Perhaps just share your thoughts around international cooperation on AI and AI development."

PAUL WONG: Oh, well, on the research sides are a lot of collaborations. Internationally Australia is well known as a collaborator, as AI collaborates all around the world. We're in good footing in that front. But in terms of the regulatory's kind of development, we tend to stand back a little bit and just wait and see what other people are doing. We have the tendencies to not be at the front of the pack, and this is quite common in relations to our approach to AI compared to what the Europeans are doing currently. They are right there at the front of the pack trying to define a framework that is workable.

BELINDA DENNETT: Great. Thanks, Paul. Look, it's been a fantastic discussion. To our speakers, we might just provide a 60 second opportunity to do a final thoughts, final wrap up before we close. Let's go to Michael for your final thoughts.

MICHAEL MILFORD: Got to go in reverse order just for equity purposes, and to give me more time.

BELINDA DENNETT: Okay. Who's ready? Let's go to Cheng. Cheng.

CHENG SOON ONG: All right. I keep harping on about data, and I think I'd just like to remind everyone again that your machine learning systems, your AI systems are built on top of the data that you collect. And my parting thought is you should think hard about what type of data you are collecting. I think you should think hard about where you should be collecting data. Maybe you can use AI to help you choose where to collect your data. But I think the quality of the data is what's important. It's worth remembering that often we collect data for a reason, and we might want to use this data for something else in the future. And if you're going to go out there as young people who are going to change the world, think hard about the kind of data you're collecting and what impacts it will have on future AI algorithms.

BELINDA DENNETT: Thanks, Cheng. Paul, some closing comments.

PAUL WONG: Yeah. I concur with Cheng completely. I want to add to that, that mix, that infrastructure is important. Without infrastructures a lot of the AI services can not go to scale. It is one thing to service 50,000 people, is another thing to service a million people. Infrastructure is important. The other important thing I want to point out is the complexities of these systems. As we are scaling up these systems, things get to be more connected. The complexity is going to increase. Once again, that questions of unintended consequence, the questions about safety of the system is paramount. And finally, last but not least, these systems are in place because we want them to provide benefit to us. Humans are at the centre of the systems, either as user or as someone benefit from the systems.

BELINDA DENNETT: Thanks, Paul. Claire, final thoughts.

CLAIRE CLARKE: Yeah. Might be a couple of things, I suppose. The first is I would just encourage people even if you don't think that you're going to be super technical and be out there building the learning models or things like that. To find out more about these kinds of things. I think having a basic understanding will really help inform future uses. It'll put you in a position to be able to say, "Oh, I think we could use a machine learning approach here. Even if that means you didn't have to go away and find someone who can do it." So yeah, I think there's a lot of value in people learning about and gaining a basic understanding of these approaches, even if they're not necessarily going to be deeply involved in the technical aspects.

And I guess the other thing I would just like to highlight is, we're sort of ... underneath this all, a human is setting the ground rules I suppose. Even if you have a machine learning system that's making decisions, that's drawing conclusions. Somewhere the very beginning is a human who's said ... my brain's gone, sorry.

BELINDA DENNETT: I think we make sense of that.

CLAIRE CLARKE: Yeah. You provide the framework in which the machine operates I suppose, is what I'm trying to say. And it's probably worth keeping that in mind. What goes in is what you've told it to look at I suppose.

BELINDA DENNETT: Thanks. And Michael, your final 30 seconds, 60 seconds we're right on time.

MICHAEL MILFORD: Sure. I guess just to echo some of the previous thoughts, we want all of society, so not just the well-educated people on this call to have an intuitive understanding of what AI can and can't do, and sort of demystify it. And so I think it's beholden on all of us, especially the future leaders in the space to make sure that we really engage with all sectors of society, because if we don't get them on board fundamentally we're going to hamstring our ability to really make those positive transformative changes in AI that we'd all like to see. I think that widespread societal engagement is going to be vital to all of the things that we've talked about today.

BELINDA DENNETT: Thanks, Michael. And on behalf of IPAA, thank you to all our speakers today. There is a small gift that I am told is award-winning handmade chocolates from local Canberran chocolatier, Jasper and Myrtle. And virtual panellists will be made, mail their gifts. Thank you to everyone for tuning in today. I hope you've enjoyed the event. On behalf of IPAA ACT, we look forward to seeing you at future event.