Institute of Public Administration Australia
ACT Division

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

ABN 24 656 727 375

TRANSCRIPT OF PROCEEDINGS

ARTIFICIAL INTELLIGENCE AND THE PUBLIC SECTOR

KEYNOTE ADDRESS

PROFESSOR GENEVIEVE BELL, DIRECTOR, 3A INSTITUTE

PANELLISTS

JAMES KAVANAGH, MICROSOFT AZURE ENGINEERING LEAD, MICROSOFT AUSTRALIA

CHARLES MCHARDIE, CHIEF INFORMATION OFFICER, DEPARTMENT OF HUMAN SERVICES

TATHAM ODDIE, MANAGING DIRECTOR, READIFY (A TELSTRA COMPANY)

ALISON LARKINS, DEPUTY SECRETARY, DEPARTMENT OF HUMAN SERVICES (CHAIR)

GANDEL HALL

NATIONL GALLERY OF AUSTRALIA, CANBERRA

TUESDAY, 20 MARCH 2018

Genevieve Bell: Oh, it's exciting to get to be here, this is a beautiful space and I'm always happy to get to come and talk to a group of people that are now closely my people again. So yay for that. I also want to acknowledge that I am back on Aboriginal country, one of the loveliest things for me about coming home is that I get to say that every time I stand on a stage. Having spent 30 years in the United States, I have to tell you that's not something they do there and it's striking to me both in its absence and now its presence here. I want to acknowledge that I'm also standing on an honourable country and that I'm happy to be home to do that.

What I wanted to do for the next 25 minutes and I've flogged someone's watch to keep time here is talk to you a little bit about artificial intelligence and how we might think about it and how we might talk about it. Now I do that conscious of a couple of things, artificial intelligence is a much bandied term, it is an often ill-defined term and it means many things to many people. I want to suggest that one of the things about artificial intelligence is that like many things, it's always just a little bit over the time horizon and that when we talk about what we mean by it, lots of people conflate everything from machine learning to deep learning to algorithms, to data, to autonomous machinery all under that label.

So what I want to do is unpack that a little bit and think about it from a slightly different point of view. Allison gave me a gracious introduction and hopped on the two things I usually do that are important here. One is that whilst I'm talking to you about artificial intelligence, I'm a cultural anthropologist both by training and predisposition. I grew up in my mother's field sites in central and northern Australia in the 1970s and 1980s and I've been around anthropologists for as long as I can remember. I actually thought that taking people's genealogies at the kitchen table was normal.

And anthropology is kind of the way that I view the world. I was lucky enough both to have that as a childhood but then to spend the last 25 years in Silicon Valley first doing my PhD at Stanford and then at Intel where I have now been, Allison scarily enough for 20 years this year. And my job at Intel was always about how did you put people into the process by which technology was built and imagined and that kind of intersection of how do you think about what people care about, what they're passionate about, what frustrates them, what they want for themselves, their kids, their communities, even their countries should be part and parcel of how we do and drive innovation.

And so that's been my job in the valley for a very long time, it's meant that I've often thought about technology from a deeply human starting point. It also means that when I came home to Australia now a year ago to join the Australian National University, I decided I was really interested in how that conversation should be taken forward. So how is it that we think about what the future of technology and our relationship to it might look. And for me in particular that's been how do we think about this coming wave of new technology and how we might locate it inside our practise as well as our communities.

I know we're going to talk about artificial intelligence this morning but my in some way starting point for the work at ANU is to say listen, AI is no more and no less than the steam engine and the conversation we're not having is the conversation about the train, the boats and the machinery. That that steam engine will ultimately power. So whilst AI preoccupies us now, it's the beginning of a much longer conversation, not the end game. And frankly whatever that conversation is and how it will proceed requires us to think profoundly differently about the world that we're building and so that's what I'm up to at the ANU. But I wanted today reflect on how we might think about the relationship between AI public policy and indeed citizenship.

I wanted to do it starting with this quote from William Gibson. I'm sure many of you have seen this quote. William Gibson is a well-known science fiction writer. He was interviewed way back in 2003 which feels like a very long time ago now about the future and he was trying to argue against a school of thinking that says the future is like this other country that someday we will move into. It will be like a nice empty house where we will just bring things, it will be clean and shiny. He was saying listen that view of the future is kind of not really helpful and the thing about the future William Gibson said back in 2003 is it's already here, it's just unevenly distributed.

By which he meant to say imagine that we should be waiting for the future probably wasn't the best idea. And for him what he was saying is if we look around we'll see bits of the future and that would drive us to a different conversation if we imagined it was here already and we just weren't looking at it. I've used this quite a lot but of late I started thinking about well, what was I looking at in 2003 that I might have also not seen as the future and maybe I should have paid attention to differently. And I realise that one of my amazing colleagues at RMIT, a professor named Larissa Hjorth sent me this photo in 2003.

She took it on a train platform in Tokyo and I'm willing to bet if you look at that photo that doesn't look like 2003, that looks like 2018. That doesn't look like that was taken 15 years ago. I mean the clothes are a bit daggy, if you pay really close attention the phones are probably Nokia candy bars, but if you squint your eyes a little bit, that could be today. That could be any train platform anywhere in the world but that was Tokyo and that was 2003. And in 2003 in Tokyo there was a phone company named DoCoMo. They had a remarkable set of services they were offering to their customers. They had early location based technologies. So they knew where you were. You could do dating by dating the people that were around you.

You could shop with things that were around you, you could use a localised mapping service. These phones didn't yet have cameras in them but if you would have looked and talked to some of these young women, they were keeping photos of their friends inside the battery cases of those phones. That was the future. We came back from Tokyo about this point in time at Intel and said to our colleagues, “Oh my God, like you should see what they're doing in Japan,” and the senior leadership at Intel went, “It’s just the Japanese, they’re weird. Also that's Tokyo, like that's just weird. Also that's DoCoMo, who cares?” And we spent a lot of time going, “That might be the future, we think this is really important,” and everyone went, “Yeah, no.”

Now imagine that the future is already around us and that's the same response we're having to it. What if there were already things going on that would help us understand what AI would look like and we're having the same response to it. We're just going, “That's a train platform in Tokyo, who cares?” So what I want to do is tell you five stories about where I think the future already is, that's already around us that has huge implications for how we might think about. The ways in which I will unfold and the questions that we should already be having. Bearing in mind again AI is a constellation of technologies, they are unevenly themselves distributed and not all of them are present. But the building blocks are already here and the building blocks start with data. You can't build some version of machinery that will predict itself and derive itself without having a steady stream of data.

If you want to think about it this way data is AI’s original sin. So what do we know about the data? Well the data is everywhere and the data is complicated. The two stories I would tell here about data, one of them starts with robotic vacuum cleaners. So everyone asks me about the robot apocalypse. If you're a Futurist one of the things people ask you is when will the world end and when are the robots killing us? Like it is a persistent question. I'm here to tell you not this week.

And possibly not before your taxes are due. So if you've got things to do, you should do them. It also turns out that the most significant installation of robots on the planet currently is the Roomba. So the robotic vacuum cleaner, a small round, sucks up dirt. Some of you may have it, it's not as common in Australia as it is in the United States. There are 10 million robotic vacuum cleaners circling on floors all over the planet, collecting dust and cat hair and crumbs. They are a fairly benign object. That is until last year when the CEO of the company that owns the Roomba said so it turns out in addition to collecting dust, we've been collecting data. We can map every single one of the homes we've ever been in, we know where your entries and exits are, those Roombas that have Cameras which is about 30% of them can tell you what the furniture is in the home and identify it by type.

We know how many doors there are in your homes, we know actually based on the dust quality who lives in your home and what kind of lifestyles you have and whether you have pets or not and the age of your kids. That should sound faintly creepy. I don't know, I'm willing to bet you think about some data that you give away and you think, “Okay, yeah I’m a bit concerned you know that.” You didn't think your vacuum cleaner was about to portray you to like the planet. Now it turns out the CEO got a sufficient pushback about this that he backed off selling this data to the nearest and highest bidders, but it made abundantly clear that a whole set of technologies we've been talking about for the last 10 years everything in the Internet of things space, every piece of technology and some of you are wearing them on your wrists, I saw them. Those of you who have fit bits and Apple watches and any number of other things.

There is an incredible traunch of data that is being collected possibly not outside with that object. But there is a tremendous wave of data that is being collected and that data is being stored and that data says something about you and that data is not necessarily yours to dispose of as you see fit.

When we think about the coming wave of artificial intelligence and its implications for public policy, one of them is about who owns our data, who gets to have access to it, under what circumstances, how long is that data good for, where is that data being kept, who is looking at it and what are they doing with it. The second one here you could just immediately kind of say and I assume everyone in the room will know is well, let's look at Facebook.

Let's think about what other data is being collected and where it is going, but the challenge here, the challenge in terms of how we think about the world we're building is about the first building block and that building block is data. The other thing you need to remember about data is that it is always retrospective which means it's always in the past. Data is about where we have been, it's not about where we are going. It can only tell you what has already happened. Second thing about data based on that same principle is that it is always partial. We don't have all the data, we just have the data we collected.

If you were at this point to decide you were going to use data to correct certain issues in public policy looking around this room, I can imagine that about 40% of us would be particularly interested in using artificial intelligence to build a better pay tool, pay parity is a public policy issue. I'm willing to imagine that 40 years on we haven't actually achieved pay parity by gender in this country or others. Say we used AI to get humans out of the equation because in theory that should lead to better pay policies. If you used all the data that existed on salaries up until now to build the tool, guess what would happen? 40% of us in the room would not be pleased, that would be the women just in case that wasn't clear because we would freeze that pay parity tool based on the data up until now which would get us somewhere between 70 to 80 cents on the dollar depending on our ecosystem and our market.

That's because the data that's collected is the world that has been, not the world we want to make, but if you want to intervene in the data to make new data sets, to change the weighting of certain kinds of data, we would need to have a very different conversation about how we are doing evidence-based policy because it would then be the counter evidence that we were needing to mobilise. So data, complicated, building block number one. Building block number two has to do with algorithms. So once you have the data, you have multiple ways that you can use data to automate processes, that's all an algorithm is. It is an automated process, it just says if this happens, then this other thing should happen.

It's relatively straightforward, you encounter algorithms all the time. This is what drives your content on Netflix, if you use it or iView from the ABC or Tinder if you are a part of that crowd or your banking system. Pretty much every time you open a digital environment, you are encountering some kind of algorithm that is automating a process for you. Algorithms work in two ways, they're built by humans who decide if this thing happens, then this other thing should happen or they're built by machines extracting patterns out of data to say if this happens, then this happens, two different sorts of algorithms. Most of the early ones were built by humans who decided if A

then B. On the one hand that seems pretty reasonable, humans know what other humans do, it should be okay.

Problem there is that humans introduce bias into the system without even realising they are doing it. The most banal and benign example I can think of has to do with a dating site in the United States that was web-based many years ago. The man who built it used psycho demographic features to match human beings and then added his own importance like a demographic feature which was that he believed men only liked women who were this much to this much shorter than them. He did not believe that men liked women that were the same height or taller and that was an impossibility to him, he did not believe it was true. This being America someone ultimately sued him for the fact that this man was never matched with the women he liked or happened to be taller than him and in the process of the lawsuit, what became clear was that the founder of this company and the writer of the algorithm had a moment in time where it was said to him some men like tall women and he went, “No.”

Really? Like are you sure about that? Because that just doesn't seem right to me. And it turned out he had built into this algorithm a world view and his world view was about how human desire worked and it never occurred to him that wasn't a human universal truth. It was his truth but he made it everyone's. So one of the challenges when humans make algorithms of course is they introduce their worldviews and they don't always know that those world views are not everyone's world views.

That's the human problem. So on the face of it that should suggest if what we did was use data to build the algorithms, we might get rid of those biases. The reality is of course other biases get introduced that go back to our problem with data. If the data is incomplete or the data is skewed in some way, the algorithm that is built on top of that data has those persistent problems. My colleagues at Google who do a lot of work in this space and who are really good about making their experiments public and their learnings public, had a really bad moment about this about eighteen months ago. When they had built a photo labelling tool that was designed to make it easier to search photos and to label photos.

And what they had done in the classic machine learning sense was they had taken a very large trove of photos, all the photos in fact in Google. They had loaded them up, they'd carved out 10%, they had run basically computing on top of that 10% to see if they could find patterns. They worked out how to label faces, they tested it on the other 90%, yeah happiness, all good. They upload this tool, first thing it does is label African-American faces as gorillas. Exactly, kind of a hash tag epic fail, not a good look. A lot of problems, Google went … And immediately fixed it. And then it became the question of how did they get there. Like how do you get to that point? Google is not a bad company, the people who built the tool weren't bad people.

Well it turns out there were two problems. One was that their entire collection of photos just tend to be the photos that had been uploaded in Google, not the photos of every face on the planet. So probably number one was their sampling method was in fact inaccurate. Because if they had

uploaded every face on the planet, they would have been a lot more black faces than there actually were in their sample. Problem number two is that the people who looked at the tool sold their own faces in it and didn't notice the absences. So here is your challenge, humans build the tool, we build in biases, we need to work out how to do it. Machines build the too, they build in biases too because we still have the same problem that the data has incompleteness to it.

Now for Google thinking about how do you manage that going forward means how do you build more diverse and robust teams. How do you make sure that the people that are looking at the algorithms reflect a whole world, not just a narrow world and how do you start thinking about the consequences. Imagine now how we think about that from a public policy point of view. How do you unbox that algorithm to know what's in it, how you working out what the data sets are, who is looking at those data sets and the rules that are being intuited there. Who is deciding what matches with what and how it is weighted.

Oh and by the way if we're building and buying or building those tools, that's one thing. We're building them on our own data. If we're buying them from somewhere else, from somewhere else on the planet, what data set is being used to train that machinery? What is sitting inside that data set that we may not like, what was sitting inside those algorithms that may not comport to our values or even our laws, and how would we regulate and scrutinise any of them. You wished you'd had more coffee don't you? I know, me too.

Problem number three if you go through these problems of kind of what does the future look like has to do with privacy and trust. This sign appears on every trash can on the West Coast of America outside of a chemist. Particular branded chemist but every chemist has a sign like this in the garbage bin. That tells you something about the world that's kind of remarkable. It tells you I mean disturbing things about the world frankly, but it tells you something about the notion of where do we sit with trust and how do we think about who knows what about us and who are they telling. And you've all had moments of this and you've wondered about it. My personal moment of this has to do with an ATM machine wishing me a happy birthday and me wondering briefly how the ATM knew that it was my birthday.

That same ATM wishing me a happy anniversary six months later was much more troubling, in no small part because I was single. And then realising that the anniversary they were wishing me was with the bank which was creepy, again. And caused me to have all sorts of questions. Bearing in mind this is a bank with whom I have a great longstanding fiduciary relationship, so the fact that they knew how long I had had a relationship with them is not surprising. The fact that they thought I wanted to know about that was deeply troubling. The challenge here is what do we know about human beings and how do we choose to tell them and how do we keep that material safe, and what does that look like.

Think about the unfolding challenges in the United States at the moment about what does it mean to have personal data, who has access to it, what

does it look like. We have new standards that rolled out in Australia just last month about breaches in data privacy, but it's not just about data privacy, it's about what are the decisions made on top of that data. We will see what happens with Cambridge Analytica as that story unfolds but what is clear there is that once they had access to the data, they only needed it once to build their models. They didn't need the data again. So one of the challenges here is if you need the data to build the model, you don't need constant access to the data to build the model, you just need a onetime access. So how do we then think about what that means. One breach is already one too many. So what does that mean in terms of how we think about public policy, how we think about decision making, what does it mean if the collision of two data sets creates a new data set that is somehow unregulated and different from the previous two, and what does that look like.

Think about that example of using data to build an algorithm, the algorithm is effectively a new world view, how is that covered, how is that scrutinised and if so by whom and what are the implications and the people you would go to, to say I don't like this world and I don't know how to fix it. Who you trust, why they are trusted and under what circumstances is an interesting puzzle. Five, nearly six years ago Mark Zuckerberg stood up in a quite public forum and declared that the privacy genie was out of the bottle and no one cared about privacy anymore. It turns out he was wrong and it turns out in fact he was wrong for himself. He now cares about privacy in a way he didn't five years ago and consumers do too, and their behaviours have shifted dramatically, both in their relationship to how they treat consumer organisations but also how they think about government. And its import to think about those as two different things.

How we operate as citizens and how we think about ourselves as consumers, we sometimes use the same language to suggest those are the same thing but the reality is they aren't. They are fundamentally different. Our notion and constitution of ourselves as citizens is very different than our relationships as consumers. Countries are one place, transnational companies are something else altogether. How you manage those things is quite complicated. What is very clear is in the last five years in the West in particular and in the US more specifically, we have seen an incredible rise in consumer awareness about where their data is moving and how it is being used and a retrenching in certain kind of practises.

So the cohort that Zuckerberg was most convinced would share everything basically millennials turn out to be deep careers about privacy and are managing their privacy both through abandoning of certain kinds of platforms and services and then embracing of ones where they can control their audience and control the media. So whether that is about the remarkable drop off in use of Facebook among people under the age of 25 … so under the age of 25 most of Facebook's audience has gone away. It turns out if you're under the age of 25, why do you want to hang out on a platform where your parents and your grandparents are. It's an age old question. Would you go to a club with your parents? No, probably not. And it turns out that cohort are using a whole range of other different platforms. Whether it is SnapChat, other forms of limited time-based sharing. Not abandoning

digital, just abandoning the broadcast aspect to it, we know that's happening.

So people are thinking profoundly differently about their privacy and it is one of the spaces where it is clear that privacy is not a linear trained with a clear endpoint of non-privacy, it is in fact a constant recursive loop. Where people are thinking more and more about who they want to share things with and under what circumstances, and the line here is that governments sit in a very different ecosystem of trust than companies do. And playing that line out isn't just a matter of saying, “Well Facebook gets to do it, maybe we should too.” So how that gets played is really different. The fourth thing in the future that's already here that we might need to think about has to do with well robots and existent artificial technology systems, autonomous systems if you want to think about it that way.

We know there was an accident yesterday in Arizona in which a driverless vehicle operating in autonomous mode although it had a driver in the seat killed a pedestrian. We don't know a lot more about it than that. What we do know is that chances are the computational system was operating in the autonomous mode. If it needed a human to intervene at that point, one of the challenges here is that computations need full human supervision, it happens faster than humans can reengage to supervise. Computation thinks very, very quickly. Its need for human supervision often demands a human engage in microseconds and humans don't work that way. We actually need more time to recalibrate into a decision. All of us know that, if you’re multitasking, we know what that looks like. We know how long it takes if someone says, “I need you to pay attention now.” It's like, “Alright, okay, good.” We can do that and it feels quick, for machines it is interminable how long it takes us.

So we know there are challenges and how we manage the trade-off between autonomous machinery and humans, we've already seen the complicated things that happen in the fast trade-off. This isn't the first accident there’s been, it's the first one we know where it's involved a pedestrian. We know there was an earlier incident in Florida 18 months ago where someone was operating their vehicle in autonomous mode and ran into a truck. Again it was a problem with both the computing vision piece of it but also asking humans to make that trade-off. One of the challenges here of course is about how we're going to manage having computational objects augmenting our decision making and our activities. For the last 250 years what automation has meant was it augmented our bodies and our physical selves.

Machines lifted things we couldn't move faster than we did, sorted things, built things. It was mostly about our physical selves that were being augmented, not our intellectual and cognitive selves. Augmenting our cognitive selves suggest something very different here. What are going to be the handoff moments between the compute object and the human? How do we think about remapping jobs, where certain tasks become computational tasks and other tasks become human tasks? What those tasks look like and how we manage the boundary line is incredibly complicated and yet in some ways be unpacked. What we do know is that where computation will work

most effectively in the short term are places where it is data rich, rule heavy, stationary and homogeneous.

Data rich, rule heavy, stationary, homogeneous. There is a reason why IBM Watson's first places are all in oncology. Lots of rules, lots of data, same object you were looking at over and over again, the learning and the knowledge is easy to imagine. Thinking about those tasks, task that historically were really privileged ones, frankly is a little bit tricky. Those are jobs that we spend a lot of time acquiring all the knowledge and the knowledge can now be supplemented and the actual pieces that the humans still need to do are the management, the framing of the jobs, the framing of the questions, the training of the machinery, that's a very different set of jobs than we're used to thinking about.

So 5th piece of the world that we still see coming that we need to think about. This sadly slightly washed out photo here was taken at Stratford on Avon two years ago when the Royal Shakespeare Company put on The Tempest and for the first time in human history we did real motion capture with The Tempest. We made aerial fly in a theatre. There's so much computation in this room you could keep the space shuttle in space, but what we did was make magic instead. And it always leaves me with this interesting question. For 200 years we've taught about the role of technology in our lives and we have measured it by efficiency and productivity. We've measured it by gains in speed, gains in the amount of stuff that could be done basically, it was an efficiency trade-off.

I wonder if that's the right metric for thinking about this future. Is artificial intelligence just about more efficiencies and if so, how would we measure human efficiencies in the first place. Certainly in the law area you might think about this as do you need less billable hours because you now have computation doing paralegal activity. That may not be the only way to think about it. What are going to be the other metrics we might want here? If not efficiency, is it efficacy because we're talking about whole systems? Is it about engagement because we're changing the model of what this work might do. Is it also about other things that we are far less used to thinking about technology giving us but the space to make new things?

Is it about creativity and God forbid the much ballyhooed word innovation? What will it be? And what if the only metrics here aren't just efficiencies and productivities, what would that open us to thinking about? And how would we imagine what the gains are here. So that leaves me before I disappear off stage and into the wilderness in Canberra, to ask you a couple of questions. Which is as you start to think about where AI is now and where it is going. We can already see the arc of the future that will unfold. This is about data and about algorithms and about machine learning and about new forms of sensing technology but it's also about ethics and regulation and reason. It's about how do we think about the systems that will contain these systems. How do we think about the ways we want to unpack and unfold all of it, and it isn't just enough to say that you know AI is about autonomous machinery because in fact there are going to be these questions about security and regulation. and transparency and explicability and trust, and all of those

need to be part and parcel of our conversations and frankly part and parcel of our regulation and our public policy.

So for me I think that raises really three kind of critical questions which I've outlined here. One of them is that last point I was making about what are the metrics we want to use as a government, as regulators, as other kinds of engaged human beings. How do we want to measure the success here? This can't just be about … it took us less time to make a decision, we might want to think about quality of decisions. There are all sorts of mechanisms we might want to involve here as we think about how would we want to use these thing. I mean I will flag here, the German government last summer, northern summer put out a set of guidelines for how they wanted to think about the introduction of autonomous vehicles in Germany. And one of the things they said was that we're not willing to contemplate having autonomous vehicles on German roads until we've hit one particular indicator of safety, and it wasn't about the drivers, it was about the pedestrians.

And their argument was that cars are part of an ecosystem in which human beings live and in fact the standard for safety there should be a whole of community standard so that it was about people outside the vehicle, not people inside the vehicle. So the mechanism there isn't about efficiency, it's about safety but it's not about the safety of the drivers, it's about the safety of the community. That's a very different way of thinking about standards that we might otherwise have proceeded with.

Second set of questions I think are really important here are about … this was put to me by someone recently who said to me, "Well where's the robot going to sit on the org chat?" Which I thought was an excellent tactical question. "Like will the robot be in charge," this person said to me. "Is the robot going to be my boss or am I going to get to boss the robot around or God do I have to work with the robot? What's that going to be like?" And it was an incredibly practical question but I think tactically it raises this interesting challenge of how is it that we want to think about including these objects in our existing workforces and how would we imagine regulating that. Will artificial intelligent objects be in some ways at humans behests, and if so how do we think about that and how do we avoid some of the language and frankly challenges that come when you imagine that objects are at our disposal and at our control, where we know in fact we need to regulate how a society treats its poorest and its most disenfranchised members of which you might argue these would suddenly become that.

If they are going to be in the pollens of the youth, our robot overlords, how do we want to think about that? How are we going to imagine that and if they are our peers and our partners, what are the organisations we would build that made that true? What would we think about in terms of how those objects sat inside our organisations both digitally and physically? How does that mean we need to think about our workplaces and our standards? All of those things become true. And then last and honestly it's because I do sit in the university and I do think about training and pedagogy a great deal in my current world. How is it that we are going to train this workforce

because oh by the way, it isn't enough to say we're going to train them in primary school because you all need employees now.

So how is it that we are going to re-imagine the mechanisms by which we do knowledge transfer and training and upscaling and that entire bundle of things if we imagine it isn't just about building another four-year degree programme. What does knowledge and training and skills acquisition look like in the 21st century when … as a like to say to my vice chancellor to frighten him, universities are the next thing that will get disrupted. So you will help that that disruption, so what does it look like and what will it be? So with all of that I want to stop, say thank you and go ha-ha to the panellists because you get to answer these questions. Thank you.

Alison Larkins: So what's going to happen next is we're going to get each panellists to provide a brief introduction of themselves and give their perspectives on progress possibilities and perils of AI and its role in the public sector. That will take us about 15 minutes and after that, we'll move to a panel discussion and then to Q&A. So we're going to start with the insider first. I'd like to invite Charles McHardie to give his five minutes.

Charles McHardie: Thanks Allison and wow, it's great to see so many people involved in the public sector that want to come along and talk about AI, ask questions about AI. So great to have you all here today. My journey with AI probably started about five years ago when I was back in defence when we were looking at how we could utilise AI to make the training experience for our joint forces, combatant forces that were heading offshore, a much better experience and I'm happy to say that's continued during my time at DHS, where we've utilised AI to find better ways to understand our large data holdings particularly in the Centrelink master programme. But more importantly to build out what we're calling our troop of digital assistants that we've utilised a variety of technologies to build out. Some of which are still in test, but some of which is in production today.

And some of the technologies we've utilised have been IBM Watson, Microsoft Cortana, and the water Microsoft bot framework, stand for technologies and also in the last six months we've started utilising Google's Tensorflow. Now has it been a brilliant journey? Now I'll talk shortly about a couple of problems along the way and a couple of opportunities. What have we built and we have in production at the moment. Our three biggest digital assistants that we have in place is Sam. Sam looks after questions that our customers could ask about Centrelink claim types in the unauthenticated space, so people have been authenticated through MyGov but they may be looking at what claims would be appropriate to their circumstance and there's a right estimator that sits in there as well to help them work out what sort of claim payment that they might be entitled to.

Once they've authenticated and they're on the inside of MyGov, we use a not digital assistant for some of our claims. Not all of them at the moment known as Oliver. And what Oliver does is help customers or future claimants work their way through the claims process. Now we're doing a lot of work in that space at the moment and come the middle of this year, we aim to release our next generation of certain claims that will be powered by Oliver.

A very different look and feel, no longer will you sit there and just go through which is pretty much a series of questions that look like an L-paper based form. This will be true conversational dialogue with a bot, a bot known as Oliver.

The third one is Roxy and we've used Roxy for staff facing. And what we've built out there is build a system where more junior processing officers when they get stuck in their claim processing task, rather than ask a more senior processing officer, they can actually ask Roxy firs. And Roy is answering approximately 80% of the questions these days that she's getting asked. The way we trained that was through a lot of interactions between junior and senior staff officers in Skype sessions. We captured all of those question and answer pairs and Roxy is doing a brilliant job these days of answering the many, many questions that our more junior processing officers have.

It's like that augmented technology that Genevieve talked about where you build out a technology that makes our staff all that more effective in their day to day tasks and that's what we've seen with Roxy. What are some of the problems we saw along the way? Particularly those at the depth sec level were very worried when we first unleashed our customer facing bots in the confidence in their answers. They said, “What if they get it wrong?” And my view was, “Well, our staff sometimes get answers wrong. If we don't start we will never get in a situation where we're able to perfect the answers that these bots can give.” So we jumped in the deep end and we've never turned back.

Performance of the understanding of these bots to answer some of the conversational dialogue that happens, the natural language classification that sits at the front of many of the sessions that customers have with these bots. You've seen it when you've utilised digital assistants. Sometimes they’ll say they don't understand or they are learning and that breaks the experience for our customers and we're aiming to put a lot of effort in making that a better capability for our bots moving forward. There are still a lot of effort that needs to be put into training and configuration of these bots. It's like training a small child, you need to continually correct and direct the bot. When you teach a small child what an apple is, if you show an apple to a child, you have to tell it several times that it's an apple and if it says it's an orange, you've got to correct it. Very similar with the bots and a lot of effort goes into that.

You've got to be careful about bias and Genevieve mentioned some of the bias. That those people that train and configure these bots need to be from a wide cross-section of society because we don't want any personal bias being introduced in these bots when they interact with customers. I think as we move further ahead with some of the bot support to decision making, we're going to come up against problems around how does some of these deep neural networks that sit behind artificial intelligence come up with their decision making process because if you give a customer a decision as to whether a claim has been awarded or not awarded at the moment, we go through a appeals process where the officer that made that decision can actually reveal why that decision was made. If you dig deeper into the AI process, how do you expose that?

So what's next for us I think better natural language classification so that you have a more seamless experience when you're dealing with our bots. It becomes more of a natural conversation and less of this sort of staccato that you get every now and then where the bot just can't answer your question. So getting a bit of flow. Bots training bots. So bots helping with that training and configuration as we move forward. Cross-session and what do I mean by cross-session? If you start a dialogue with a bot on a mobile phone and you have to end it because you need to take a phone call or jump on a train or you've arrived home, then you can pick that session up on another device.

You know whether that's a laptop or your desktop. A water bot framework, so allowing the bots to be able to reach out to other platforms like messaging platforms. So the bots if they need to push a message out to a customer, they can push out messages through push notifications to things such as WhatsApp or Facebook Messenger. We see that as part of the future and I think one of the big changes particularly in federal government service delivery is multi-lingual and better assistive technologies. So leveraging some of these platforms such as Google Translate, some of the technologies that Microsoft have, to be able to offer up those bot interactions in many different languages to reduce some of the load that we have on our translation services that many of our federal departments offer up.

And better assistive technology. So our bots being able to talk to you and listen to you rather than you have to type and look at a screen. So I think there's some of the things that we've been doing in DHS, all very exciting and we're very keen to help others across federal government or other sectors in sharing our experiences.

Alison Larkins: Thanks Charles, that's fantastic to get an overview of you're such a big shot of what's happening in DHS, and of course you're an agency that's so critical in terms of community facing, so that's great. Let's move to Microsoft, you actually set that up very nicely. Charles set it up very nicely for you. James.

James Kavanagh: Yeah, sure I thought what I might do is give a little bit of an explanation of some of the ingredients AI and particularly why it's so important right now, what we're facing. What's interesting is that we're talking a lot about the emergence of AI as it's something of a spring, an AI spring but that spring is following winter and we have had multiple AI winters. It's not a new concept the idea of artificial intelligence. And it's important to understand why some of those failed and what's different about now and why is it so important, so emergent. And particularly why is it so interesting right now to us in this room.

If you think about those winters that have happened. The first real emergence of AI was the idea of Connectionism. That we would just connect lots of systems together and magically intelligence would appear. And that was a dominant thinking in the 60s and 70s and it failed miserably. The next major wave of AI was the idea of expert systems that if we could just write down and codify all the rules of the world and all of the things in the world, if we could just do that simple task, then computer scientists would be able to emerge with intelligence, and that failed.

But now we're in the domain of learning systems and it's quite different now. It's quite different in terms of what we've been able to achieve, how much progress in the last few years has happened and how the trajectory is actually continuing to grow faster and faster. I think there are five big ingredients that support that emergence of AI and that spring at the moment. And I want to kind of go through those five and talk about why they're relevant right now to us here.

The first is data and Genevieve talked about this quite a lot. That data is the fuel of AI. It's the rocket fuel that allows us to do quite amazing things. Over the last few years from … if I just talk from a Microsoft perspective three years ago we beat the human ability to translate speech across multiple languages. That was three years ago, two years ago we're changing the algorithms and the data that we had, we were able to beat the classification of an image and achieve superior performance over human. Those kind of progress only comes from data and the challenge is trying to make available of that data so that it can be used as a training model, used for particular outcomes that government or enterprise or others want to achieve and do that in a secure and safe way.

And here in Canberra there are petabytes, exabytes of data available. Whether it's geospatial data in environment, whether it's archives, whether it's tax, whether it's human service information. Whether it's defence logistics, whether it's space satellite, this time is full of data and the challenge is securely on locking that data, making it available for us to train models and for yourselves to be able to train and deliver better service. Just like Charles talked about using the data within Human Services to actually train a model. The data is the ingredient number one, ingredient number two is having the computing to be able to process that data. There's no value to the data if it is locked away. What we really need to be able to do is to process that data with very large scale computing.

And the cloud computing platforms that exist right now deliver that possibility. What we're very fortunate and I'm very happy about is that we are building the global cloud platforms in Canberra. Microsoft 2.0 that as a cloud platform into data centres here in Canberra to sit alongside the data of government, the data of critical infrastructure to manage that security and allow the processing. And it's quite an elegant thing that if you can bring data and compute together, you can unlock lots of insights, you can share that data securely and you can drive forward on AI.

The third big ingredient is algorithms. Being able to process that data and the challenge in the past is these algorithms have not … they've not been revolutionised. We haven't developed brand new algorithms in the last three years, the last five years what we have done is refined the algorithms of neural networks, the algorithms of statistical regression and classification and kernel learning methods that have been very well established. But AI now refined in their practise and more importantly are made available for people to use. So it's possible for a software developer, a business user to actually take advantage of reinforcement learning or neural networks in a way that they might previously have required a PhD in computer science to be able to understand. So that access those algorithms is really important.

The 4th ingredient that I think is really essential is skills and those skills might be in computer science, being able to deeply understand the way AI works. It might be in applying that into practise as a developer, as a programmer. It may be in the policy of the business management space of understanding how to use AI to achieve policy or particular outcomes. A little known known fact I mean three years ago AlphaGo was a system that beat the world champion Go player in a game and really was quite a milestone. It was one of the most complex games that could be beaten by an AI and, when that happened everybody kind of stepped back and remarked in the AI world that we seem to have gone past a turning point, a watershed moment.

One of the co-creators of the AlphaGo system was a fellow by the name of Dr. Shane Leg. He studied in ANU, he did his PhD here. He's a Canberran who went to London and built that system. We have some of the best researchers in the ANU, in Uni of Sydney, in UNSW. We have some of the best capabilities with national computer infrastructure here in Canberra. And what it really is about for us is to harness those skills and those capabilities to take advantage of the data in the compute that we're bringing. And then the final one and Genevieve touched on this as well is really the final ingredient to me is imagination. And having the imagination to apply AI, to augment people's abilities to deliver policy or programme outcomes.

To give you another example here in Canberra, a little small company called Gravity Consulting, they build these most elegant dashboards for programme management and a number of agencies and departments here use their services. What they do is they surface up the status of different projects, the reports, the budgets, the risk models, the verbatim of status updates, they're showed in a beautiful dashboard. But if you talk to the people in Gravity they don't describe their role as building dashboards for programmes. What they describe their role as is using AI to understand the true state of programmes and guide more effective management decisions.

So what's actually happening within their system is that all of the data, all of the reports, all of their verbatim, the emails, the risk reports, registers, everything relating to a programme is analysed, processed and interpreted by machine learning algorithms. And it can be quite a striking thing to see programmes that are showing up as bright green, all on track all delivery looking like it's going to be great. But the algorithm is saying this is bright red, it's burning and pretty soon it's going to go on fire. And that's applying artificial intelligence to quite an important kind of programme objective.

I really think that we're now in the space of moving into a summer of AI where it's really broadly going to be applied and I think there's some unique characteristics of Canberra in terms of the policy ability that we have, the data that exists here, the compute that's here, the skills that exist in this particular market. But this is a great place to really advance AI and it's going to be tremendous to see how that plays out. Thank you.

Alison Larkins: Thank you James I think you've raised a number of really interesting issues that we might investigate more in questions. The issue about data given the challenges that Genevieve is raising and the issues around trust. Which I'm sure Charles will want to talk about as well. But also I don't know is as a

programme manager myself I'm totally terrified by the idea that there might be a different source of information about how the programme is going. So it will be good to talk about that in a bit more detail. But let's move to Tatham now.

Tatham Oddie: How does it feel to go fourth after all these wonderful people? I think on the progress first of all, it's amazing that we can have 300 people in a room like this talking about something which is really a whole bunch of nerdy math under the covers. And looking at how do we go and apply that to what we do but also at a very early stage of that journey and that we still call artificial. Computers used to be humans who computed things and worked out maths on paper and we don't call the current ones artificial computers, they're just computers and the progress in that change is the fascinating film of hidden figures. So it will be interesting to say when does the word artificial actually drop out of the phrase and the terminology that we use.

I think in terms of the potential we've had some great stories here and the particular one that I really like to focus on is around kind of flipping the idea of efficiency to how can we expand the group of people or customers or citizens or whoever it is that we're helping. So really in that kind of scale piece where it's say we're going in processing claims, how can we get through a long backlog of claims that historically would fold to the bottom of some pile and never quite get there or for offering legal advice, how can we go and get to a larger community. In the case of say the cancer and the oncology scenario, how can we not just be more efficient in the first world but how can we take that type of processing to third world environments where they might not have access to the same set of skills.

So how do we rapidly expand through augmentation of these artificial models, our ability to go and scale this out. But they all sound like really big complex things to implement sometimes. So the story of you've got three excellent bots out in progress in DHS alone, there are a really big complex thing to deliver. I also like the idea of where does kind of small AI pop up? Where does it pop up in the little things we're doing each day and one example that I love is the latest version of PowerPoint, really simple. You make these slides in it, you drop a photo into that and it now automatically goes and actually adds a caption to that photo for accessibility of it. It doesn't always work terribly well.

I used it at a talk a couple of months ago, put in a photo of a Syrian refugee camp and it came out a group of people in a beach. I don't think they quite felt that way themselves, it did also flag it as auto-generated and I had the opportunity to go and correct that but that's going and taking content that I would produce which historically might not be very accessible and then making that more accessible to a larger audience. We go and combine these building blocks of things like real time translation, image analytics, natural language processing and then we start to be able to scale that to incredibly broad communities.

The peril through all of this of course is how do we trust it. It's come up a few times, how do we diagnose the decisions that have been made. We're right on the cusp of just daily, we'll actually be able produce some useful

predictions or analytics or decisions but we don't understand why they're occurring. And the kind of best example I see at the moment of how we might be starting to solve that is actually around GPSs, in terms of the navigation systems we use in our cars. We've had them 10, 15 years now and it starts with people saying, "I don't trust it, that's stupid TomTom. Why is it telling me to go that way?" And you don't understand what it's trying to wrap you around. You don't know that there's an accident down there or that there's some traffic or that the roads changed since last time you drove there, and you've got your own preconceived idea.

It has a name in TomTom and you need to start to trust that name. Now there's multiple choices, you've got Google Maps and Apple and Ways and all these different types of algorithms and they each have a different style to them. So Ways is a very aggressive routing algorithm. It will turn you a corner before the traffic lights so that you can turn again so you can skip the lights. Some people like that, some don't. So you get into an Uber these days and you can actually ask the driver and say, "Hey, do you mind if you use Ways?" "Yeah, sure I like Ways. Or no, I usually prefer to use Apple Maps," or something like that. So you start to get this identity, it's attached to them and then a diversity of algorithms that can solve the same problem in different ways.

As we think about cohorts of bots and diversity of the teams who prepare these types of algorithms and train them, how do we have a diversity of bots? So I'm challenging your scenario where we've got something that says, "Support stuff," would the stuff have a choice of maybe three different solutions they can go to which give them potentially three slightly different forms of analysis in the same way that they could go to three different senior staffers for advice and then go and compare those results. So trust and understanding is going to be one of the really interesting things for us to solve still.

Alison Larkins: Okay, thank you, that's fantastic. And that's started us really on the first question that Genevieve asked what metrics do we, how do we measure success. What metrics do we need to bring to the table as we start applying AI? And I think you're arguing for social good diversity abroad, a sense than efficiency and-

Tatham Oddie: I think it's essentially a form of efficiency but started in the form of whatever your current customer service or success metrics are. So if you have a current set of customers, how do you help more and help them in a better way. Now if that means that you need to be more efficient and the algorithm can help you do that, spectacular. But fundamentally we don't have metrics of how many computers we use in your organisation, we have our goal for our respective organisations and business units and those are the goals we need to stick to. We're just getting better at delivering those.

Alison Larkins: Charles do you have a comment on that first question, on Genevieve's?

Charles McHardie: Yeah, so the way we generally with our digital assistants has worked okay but I think we need to move beyond this is the like, dislike and half like. And we've used that as a way of training and configuring the bots. So if we get a

like, it generally means you've answered the question correctly, so the bot knows it's good on that sort of question type. Half like means it needs to go back and reassess and if it gets a thumbs down, then we put manual intervention from one of our training configurement offices to be able to assist it because the bot's struggling in that area. And the other metric we tend to use is just how many questions have been successfully answered. So percentage of questions being correctly given a thumbs up.

Alison Larkins: So that is in a type of efficiency, but usability is pretty critical.

Charles McHardie: It is yeah.

Alison Larkins: James do you have any comments on Genevieve's challenge?

James Kavanagh: Yeah, I think it's an interesting very human way of thinking of the problem.

Alison Larkins: Well, that's a good thing, isn't it?

James Kavanagh: Well, maybe not. So let me put it this way. As humans we tend to look at metrics as quite reductionist way. We try to look at one number as being important like GDP growth or your pulse is an indicator of health. But machine learning algorithms don't work like that. You put in tens of thousands of parameters and what the algorithm is trying to do is to optimise based on a very, very wide set of parameters. You don't tell a machine learning programme here's the key metric, optimise only on that metric, what it's trying to do is to solve across massive scale of parameters. So the interesting thing is if you say well, it's about metrics. Well that's a human way of thinking of the problem. If a computer is trying to optimise a decision process of assessing data deciding and acting, it's not optimising based on an individual metric, it's optimising on tens of thousands of metrics and trying to maximise utility function that comprises all of those.

So first of all, at the human way of thinking so we have to kind of maybe think about it a little bit differently. I think one of the key metrics though from a human perspective is important is resilience and adaptability. With so much change happening and us expecting automation in AI to perform a role within our organisations that are constantly changing, constantly having to deal with disruptions, threats, attacks, new opportunities, somehow we have to think about how AI will help us to be more resilient, to be more adaptable and to find some kind of metric or guidance around that. Traditional metrics of resilience are very technical. They're like time to recover from failure, availably levels. They are not business, they're not social resilience measures and somehow I think that that's the metric that we need to try and think about.

Charles McHardie: And I might just say, you can tune these platforms as well. If you look at the three core steps that a bot will take, it will do the natural language classification peace up front where it tries to understand the context of your question. Then it will move into the deep neural network and to come up with a range of answers and then it will go into a probabilistic determination to work out of the list of answers that's come out of the deep neural network, what order of probability, which one is the correct answer. Now

you can actually tune it up so that you will get more answers but you've also got more chance of getting wrong answers or you can go the other way where you'll have less chance of having wrong answers but you might not get an answer at all. So you tread that sort of thin line as you're trying to tune the bots based on the feedback you're getting from your customers.

Alison Larkins: So what about Genevieve's question about where these technologies sit on our organisational chart. Have any of you got any views on that?

Charles McHardie: So I might approach it from a public service perspective. I think it makes sense to give these bots a persona for a start. Whether they are external facing or facing within the department because staff tend to get some buy in because they are helping, trying and configure these bots. So I think that's important. Where do you sit them in the org chart is a tough one because they are basically multi-discipline teams that are looking after these bots both technologists but also the business folk within the department that are looking after it. I don't think they sort of sit quite neatly within you know ICT or within one of the business areas.

Alison Larkins: So they're integrated in in fact into the workforce.

Charles McHardie: They are.

James Kavanagh: Yeah I think that's pretty a tough question. Most of these technologies right now are augmenting people's ability. So they're actually sitting alongside, they're an effective part of that role. I mean some experiences though that we have is … we use Skype a lot. Inside we're always kind of instant messaging to one another within work. And often when you look for a question or you're looking for somebody a bot will appear as one of the choices you can ask the questions to. It's not always ideal, it's not always the right answer. Often it can just lead you a part of the journey but they appear as people that we can contact and we can ask questions of. I think increasingly that will be the case that it's not so much about where they sit in the organisational structure but it's more about the functions and how we answer our problems, how we solve them, within that organisation.

Alison Larkins: So are any of you actively engaging with the question about skills and training that Genevieve posed as her last question. I think she's basically saying we've relied on a three or four-year degree, technical degree to get you into a career working in the field. Are any of you now starting to disrupt that model or thinking about disrupting that model in your own organisations?

Tatham Oddie: I'd say it's a model that's never fundamentally worked for us. We're a 17-year old technology company and what we do now is very different to what we did last year or 17 years ago. There are a set of base principles of ways of thinking that we get out of that more kind of traditional training type approach but also we have less than maybe a third of our organisation I think, are actually degree-qualified. We put a lot of effort into it as an organisation our capability development which starts from how do we identify what are the capabilities and the skill sets we need, what's the trajectory that they're on? So we have a formal methodology for going and

doing this, making sure we then communicate that to our stuff. Where do we want to invest in skills and capabilities but also actually consciously pointing out to people where do we want to de-invest in things and start to scale them down over time because we see them being less relevant.

We then see they need to have a commitment to well how do we give people a bridge to what's next. So even with our data analysts, the types of work that they're doing is changing and we've had a bit of a focus, move away from say econometrics into more natural language processing. So on our capability and our skills map that we regularly go and update. So regularly for us is we do a major update twice a year and then a minor update each quarter. We actually go and document out and say if you've got skills in this area, we think you should probably go and look at this area next.

We don't go and then dictate to people where they need to go, maybe it's time that they roll out of the organisation and that's okay. But make sure that we demonstrate a clear path constantly of where do we think is next.

James Kavanagh: So you could probably imagine in Microsoft, an engineering perspective got a very technical workforce. And most of them are exposed to a lot of the more advanced technologies that can work in this space. But we found it particularly difficult to retrain and extend some of our engineers into this space. So we actually set up a AI university inside Microsoft. It's pretty significant run by our research teams. They do programmes that last about six months that any engineer can effectively enrol into and be trained over the course of six months. And leads to a kind of a capstone project that effectively gives them an internal professional qualification in AI.

We've since explained extended that outward. So we now deliver the same to edX and Coxeter and another programme. So our customers and others can take advantage of it too. So for us it's been a challenge of you can't assume that an engineer who is very capable technically will be able to understand and apply the maths and the computer science and the broader concepts of AI. So we've had to invest quite deeply in building that capability amongst them.

Charles McHardie: From a public service perspective, it's been great to see the platforms evolve and how much more simple it is now to be able to interface with things like the bot framework itself. Some of the AI that now sits on top of data analytics platforms, we don't need to have any-

Alison Larkins: So you getting … you're actually changing your capability-

Charles McHardie: it's just a series of windows that you can work through. We found that it's quite easy for our folks to pick it up. I think the message for the public service is grab some data, start putting some open source AI platforms over the top and start experimenting with it and build up some innovation teams to start looking at AI. Do some rating, just get a basic knowledge around what I talked about on natural language classification, deep neural networks, probabilistic determination etcetera and people will start to pick up the gist of what sits behind it. But the interfaces that are there now are so much better than they were several years ago where you basically had to hand

over what you wanted to an engineer and let them crack on with it, it's not like that anymore.

Alison Larkins: Well that's a challenge and encouragement to us all, I think to go out. So let's now open it up to questions from the floor and I know many of you will have thought of great questions through the morning, up the back.

Question 1: This is for the whole panel. What do you think of China's Sesame score and the fact that it's built on an algorithm solely run by the Alibaba Group. Can you say anything why they're coming perhaps by stealth into our system and how would you guide against that?

James Kavanagh: I know a little about it, it's quite scary. The concept of it is if I'm right in what you're referring to, it's the Social Credit Score concept within China and the idea that based on your interactions either if you're blogging against the government, if you fail on your debts, if you engage in activities that are not regarded as consistent with the policies of the government, then your credit score can be declined and it's kind of one hit and you're out. You're forever untrusted after you've had a consequence. It's pretty concerning. That's a concept, an idea that effectively is mining data, associating data, using that for social consequence and social control in a way.

From our perspective and if I took from a Microsoft view, our view is that that is not using data effectively to drive good positive outcomes. We think it's really important that data must be in the control of the consumer, data must be always … sort of previously must always be protected and so I think it certainly seems like a concerning development that some countries around the world or even some enterprises are focusing on using data in a way that's not for the best interests of their citizens.

Would it work in Australia? Hopefully we have the right safeguards from a government, from a policy, from an enterprise and data perspective to prevent against that, but it is a concerning development that that we see in terms of how data can be used and abused for social control or social influence. That's about as deep a view as I have on it right now.

Tatham Oddie: I think one of the relevant areas would be the European GDPR legislation. One of the elements of that is the right to explain or the right to an explanation of how your data can be used and that comes back to the traceability of as a consumer if I contribute this data to your platform, you have a right to ask now what impacts could that have. So if I tell you my date of birth, you need to come back and tell me that that can compact my score in some way, that you generate as an organisation. So that's one of the rights code of fighting that legislation, and then the other is a right to delete.

Now the challenge with that though is that if you look at a platform like Facebook for example, do you have a right to delete and close your account? Sure, but there's 2.2 billion people. For a lot of those people, that is the Internet. So good luck existing outside of that. So certainly the policy element, very important in supporting that.

Question 2: My name is Wicker and I'm from the Department of finance. I'm interested to know, Charles you mentioned that we should be out there actually investigating these sort of stuff, what's your view on having a consolidated approach across government to actually investigating this stuff? Just where there are common functions and there can some investigation that's leveraged by a range of different agencies.

Charles McHardie: Yeah, so that's a very good point and indeed it's one of the things that's considered at OGAS Group known as the D10 where CDOs, CTOs, CIOs come together across federal government. It's facilitated by the DTA and indeed just at the last D10, one of our folks who's out in the audience there somewhere who is one of the National managers leading AI within CO Group within DHS came along and presented some of what I talked about today. Now that's generated a bunch of dialogue amongst the departments because several departments have different platforms like the ITO runs Alex to help people out on the tax websites and to be able to run things like MyTax.

So we are looking at ways now how we can leverage across federal government. So for those departments that have not started on this AI journey, how can they talk to some of the larger departments and leverage off what we've done already and indeed for us to assist where we can. So that that's kicked off.

Question 3: My name is Andy and I'm from the Department of finance. My question are defining the efficiencies AI algorithms and how I don't remember your name but second panel member was saying that this is a very humanly way of thinking about it. Can you a bit elaborate on what you meant by that and what are the alternatives because if we can define those matrices and how we define efficiencies and even like the example that Genevieve was giving about safety of sort of self-driving cars, how can be control the outcomes because of course the AI algorithm is going to have an accuracy. Are there measures that we have but how can we control that outcome?

James Kavanagh: it's James. So let me try and explain this. So if you're trying to train a machine learning algorithm, the way in which you do that is provide very large sets of data that have many, many parameters. So you might let's say you wanted to figure out well how can we model house prices within all of Australia. You wouldn't say these are the four or five most important things, your location, the number of bedrooms, whether you've got a garage or not for example, you wouldn't set those parameters. What you give is all of the data you can possibly provide. Including pictures of the house, including the inhabitants, the location, the address, whether it's got a one in the address or it doesn't or an H and therefore it's lucky or not. You just provide massive amounts of data and you train the model not by saying this is the most important metric, this is the most important parameter. You train the model and allow the computer to decide what are the most important statistical parameters.

So that's why I say what you do from a machine learning perspective is not like what we do from a human perspective. We think in terms of quite productionist, there're set of key parameters and we think in that way. So

that's what I was trying to relate to is metrics are important to us, a machine learning algorithm doesn't think in terms of metrics. What it's always trying to do is optimise a utility function, you can think of that utility function as happiness or you could think of it as outcome whatever it is, but it tries to optimise that utility function based on a massive number of parameters. So it doesn't try to reduce to a certain set of metrics.

Question 3 (follow up): So thinking like that I think to me it's a bit scary and I relate to that, because I'm coming from AI background. So say like a self-driving car, the decision is to get to point A to B through the fastest road and in your way there is a pedestrian. So if our focus remains on getting the most optimal outcome, it's sort of a scary thing. I think us as people who are in control of those algorithms can change the area of focus as well. So of course the algorithm is going to give us the most like sort of the best performing outcome that we ask of it, that's quite natural. But we can change the area of focus and I think like with the bus that Charles was talking about with the sort of, that they have to give the best answers to the customers of course but what about the empathy? Like how can we sort of make them more personable? Those are the things that we have to be conscious of, because without those things I think it could be some scary futures and the movies that we see that can come true. So that's what I wanted to add. Thank you.

Alison Larkins: Now there was a question down here, yes.

Question 4: Thank you for your presentation, I found it very interesting. I'm Louise from UK from the office of Parliamentary Counsel and we draught legislation for the Federal Parliament, and my interest is in AI and the application of the law and in particular legislation. My question is do you think that in the future bots will be able to provide legal advice? And I was particularly interested in Charles what you were saying about the three bots that you have working for you where, I think you said that they ask a series of questions and I can see that for particular provisions like particularly ones like Social Security or family payments where it might be eligibility and there might be a set of criteria, and it's a yes or no answer, and I can see that bots would be able to that by being programmed.

But in the future for ones that aren't a yes or no answer, do you think they will be able to handle that? And related to that question is, do you think that there might be changes in the way that we draught legislation to accommodate that and third the difference between automated decision making and AI.

Charles McHardie: Wow, that's going to cover a bit of ground but that's okay. So the first one, first question the answer is yes. They are now dealing with questions that are more complex than yes or no. It's the way you've asked the question and you've got to play it back. So if you look at settling claim, the bot will collect the information, you'll get to a customer's summary at the end and it will play that back to the customer and then you have an opportunity to correct anything you may have given to the bot. Very similar to when you're putting in MyTax form, you get to the end and there's a summary. So that can be done, that's no problems at all.

The legislation piece is an interesting one, I remember when we first started looking at utilising bots and our chief digital officer at the time, Tam Shepard said to me, he said, “Charles this thing's rubbish,” he said, “I thought all we were going to do is get the operational blueprint which is all the policies and procedures for awarding claims, give it to it, let it read it and then it will become the expert.” And I said, “No, no it’s not that at all.” And I think James made mention of that. There is a lot of training configuration that needs to happen behind the scenes. You cannot just give one of these bots a whole pile of legislation and it sieves through it and becomes the expert.

Now you can use AI to help in your understanding of that legislation indeed when there is massive volumes of it, and I think IBM have done a very good job and Genevieve mentioned it, start with Watson around things like oncology etcetera, where doctors need to read a lot and they just don't have the time to do it. Watson can help them do that. But you can't just give it a massive amount of data and expect it to become an expert overnight, you've got to train and configure those platforms. So I think we've got some way to go before the lawyers can use it quite easily. And I know some law firms are using it to help with better understanding of quite complex legal policies, etcetera.

Decision making is an interesting one and if you look at it from a Centerlink perspective, would we get to a situation where a bot could make a final decision on awarding or rejecting a claim. Now I think that’s someway off but I think you can get to a point because there‘re large insurance companies in the US now that are using AI to turn around automobile insurance claims quite quickly. Where you have an accident, you use your smartphone, you take some pictures of the accident if there's a police report, you upload it and they’re using AI to actually come up with a decision in like 40 minutes. So there’re other industries doing that but I think Genevieve also made a very good point that just because you can do it in the commercial world, doesn't mean you can do it in the public sector. So I think we've got some way to go.

James Kavanagh: If I could add one example where it's already happening. So there is a book called Do Not Pay in the UK and the US right now and Do Not Pay basically deals with parking infringements and fines. So if you go into the Do Not Pay in the jack bot, you put in the details of your fine. It’ll figure out whether you have grounds for appeal or to delay the fine or whatever. So far it's processed about 500,000 cases and in about 180,000 cases been successful in overturning the fine. So it's already happening in a way.

Tatham Oddie: You'll need the bots to keep up against it. Some of the types of information that can come out of this topic of analytics is not just decisions but also suggesting. Have you considered X, Y and Z. Maybe this legislation is relevant and it's somewhere in the back of the brain or you just don't know about it and those types of things. So back in that kind of supportive model as well, where it doesn't necessarily need to make a decision but just give you some hints of where to progress next, in particularly complex areas. In terms of the automated decision making thread, some of those kind of three stages we think about in analytics is first of all descriptive analytics. So if we look at

something happened, well why did it happen? So something that can go on tell us why.

Then predictive analytics, what might happen in the future and then the third stage is we build up enough confidences that in that is then prescriptive analytics. So based on all of these inputs, actually prescribe the outcome and go and make the decision. So that becomes a bit of a kind of maturity steps that we can work through to build confidence in something over time.

Alison Larkins: Look, I'm sorry we've run out of time and I can see that there's still a lot of appetite for questions but I hope you will all join with me in thanking the panel. It's been a really fantastic conversation.