Artificial Intelligence: Can Australia Chart a Different Course on AI?

Ed Santow, Lizzie O'Shea, Peter Lewis

Ebony Bennett 0:02

I have to wait for a little announcement from Zoom telling everyone that recording is in progress now before I kick off. Gidday everyone, I'm Ebony Bennett. I'm Deputy Director at the Australia Institute and welcome to our webinar series. Thanks so much for joining us today. I'd like to begin by acknowledging that I live in work on Ngunnawal country here in Canberra, and pay my respects to elders past and present and I also want to send a big congratulations to the Uluru Statement from the Heart, which just this week was announced as the winner of the Sydney Peace Prize. And I hope that we'll all redouble our efforts to work towards a Voice to Parliament Treaty and truth telling in this nation. The Australia Institute does these webinars, hopefully weekly. Sometimes we've got a bit of a gap as we do next week. But you can find all the details for these events at australiainstitute.org.au. And just a few tips before we begin to help things run smoothly today. If you hover over the bottom of your Zoom screen, you should be able to see a Q&A function where you can ask questions of our panelists or type them in. And you should also be able to upvote questions from other people and make comments. We've got a chat going - please keep things civil and on topic in there or will boot you out. And finally, even though you get that reminder from Zoom, a reminder that this discussion is being recorded, and will go up on our YouTube channel afterwards, if you have to duck out for any reason. So we've increasingly complex algorithms making ever more significant decisions over our lives. There are several really important questions that we need to start considering in Australia. Like should governments deploy facial recognition technology? Can decisions that affect people's rights be automated? Or should they? Should algorithms be forced to comply with Anti-Discrimination Law? And all these questions are more really big and complex. New technologies like AI are reshaping our world and they bring powerful opportunities and threats, especially to our human rights. And a new report from the Human Rights Commission sets out a roadmap for Australia to seize the opportunities and address those threats. And so today, I'm delighted to be joined by three guests who are going to shed some light on some of these big questions, and how tech will affect us going into the future. So please, let me introduce you to Ed Santow who's been the Human Rights Commissioner at the Australian Human Rights Commission since August of 2016. Ed leads the Commission's work on technology and human rights, refugees and migration, human rights issues affecting LGBTI Australians and counterterrorism and national security. He is also a fellow at the Australian Academy of Law, a visiting Professorial Fellow at the University of New South Wales, and a member of the World Economic Forum's Global Future Council on Human Rights. I'm also joined by Pete Lewis, the Director of the Australian Institute's Centre for Responsible Technology, and Executive Director at Essential Media. And last but not least, I'm delighted to welcome Lizzie O'Shea, Chair of Digital Rights Watch. A big welcome to all three of you. Thank you so much for joining us. And Pete, I'll hand over to you.

Peter Lewis 3:11

Hi, thanks, Ebony. And hi everyone. Ed and Lizzie, great to have you on board for this discussion today. I think it is a really significant day in Australia. This is one of those reports that is the result of years of work and consultation and what's come out the other end is not just nationally significant, but globally significant. So hopefully in this hour we can give you will a bit of a taste of what's being put in front of the government to lead in the thinking of the way technology develops, but also critically, to think about how to position Australia not just as a follower on tech or someone that puts out bushfires, but someone that really leads on the way that tech develops in the future. One of my sort of analogies I used for Ed's report is that the way that the Swiss make really good watches and the Danes make really good furniture, can Australia make really good AI technology? So that's our challenge to be the Danes of AI. So before we dip in, Ed, congratulations, I know you're doing the rounds in Canberra today. Do you want to give us a bit of an overview of what led you to this report? And then I'll bring Lizzie in just because she's been working in this area for a long, long time. And you will have to hit unmute or you're buying the drinks!

Lizzie O'Shea 4:36

He can't unmute!.

Peter Lewis 4:39

Lizzie, why don't you set the scene as Ed works out how to unmute. You guys have been working in this space for a long time so this is quite exciting to see something like this work its way through the process, isn't it?

Lizzie O'Shea 4:50

Think he is back, so the right to unmute is probably one of the key rights that need to be, needs to be advocated for.

Ed Santow 4:57

Yeah, I can. Yeah, freedom of expression. It's such an important part of all of our lives. Look in the intro there Pete, and I'm coming to you from Gadigal land and I pay tribute to the owners of that land, which has never been ceded. In your intro, you mentioned this has been a work of a few years. That's true. And we started back in 2018, this project, because we were living in the middle of the fourth industrial revolution. And it's new technology that's driving that change. And we were beset by these competing visions. You have the technology enthusiasts who paint this positive picture of what the revolution is bringing to us, the end of dirty and dangerous jobs and endless production line of exciting whiz-bang gadgets, Artificial Intelligence driving better and smarter decisions. But as I say, that's not the only vision. There's a competing vision. And that's, people, well, I mean, people find meaning, not to mention the means to put a roof over their heads and food on the table in their jobs. And the idea that your own job will be automated out of existence is, can be terrifying. And similarly, AI doesn't always generate better decisions. It can result in new forms of unfairness, discrimination, and worse. So in one sense, what we had to do is we had to grapple with those

competing visions and the truth is that there's some reality in both, that there is an enormous amount of promise that new technology like AI brings, but that promise will not be fully realised, unless we address the risks and threats of harm to our human rights. And so the starting point for us was to first understand better what the Human Rights implications of AI are. And secondly, to make sure that we have a clear plan to address it. And that's essentially what our final report sets out.

Peter Lewis 6:59

Just take us through that process. So you start off with "everyone's got a view". How, what's been the process of distilling it into this final set of recommendations?

Ed Santow 7:11

So we've done two major rounds of consultation. And as with any consultation, we've listened to everyone who will speak with us. So that's experts, people who, you know, have an interest in the area, potentially commercially, or others. So we've heard a lot from the corporate world and from government, from academia. But we've really tried to engage, especially with civil society organizations, and just people in the community. Because part of this process that we've been on, to better understand what AI really means, is one that others are on the same journey. And so we've been trying to share what we've learned as we've gone. But we've relied very heavily on people like Lizzie O'Shea, who is truly one of the leading experts in civil society on this anywhere in the world, to better understand the issues, and particularly the risks. And then what we've tried to do is distill it down to understand where we need to really focus attention most.

Peter Lewis 8:23

Lizzie, that's a pretty good intro from Ed to you.

Lizzy O'Shea 8:26

Very generous endorsement thank you Ed. But yeah, I am personally very excited about this report for a few reasons. I mean, the first one is I think you're right Pete to talk about the different ways in which technology, and Ed, to talk about the way in which the technological revolution has been presented to us. And one of our big concerns is that technology's often treated as a more neutral or better version of human decision making. And I think we need to work hard to criticize that idea. And also make sure there are guardrails and safeguards in place, so that that assumption doesn't conceal discriminatory conduct or unfair outcomes for everyday people. And the other part of it, of course, as many people in the room might know, we don't have human rights enshrined in law in Australia. We don't have a Charter of Rights that's enforceable for everyday people at a federal level. And we do stand as outliers and as a Western democracy, as a liberal democracy in that way. So it is really critical I think that, you know, people like Ed and institutions like the Human Rights Commission, do this work to figure out ways in which we can make sure human rights are given weight in the development of this technology, in the use of it by government in particular. So then instead of it becoming just a way in which companies can make money and governments can, perhaps, organize and trim services that they provide, without reference to the individual and their

rights, that's kind of... we don't just look at it through the lens of competition or promoting industry or promoting efficiency in government services, we actually look at the rights of the individual. So I think people like Ed have to work extra hard in Australia as opposed to other countries. But I do think that means we get much more interesting recommendations and a thoughtful process that, you know, is from the bottom up assessing how we can better respect the rights of individuals in this rapidly developing space.

Peter Lewis 10:23

Ed, Lizzie's talking there about guardrails and red lines, and maybe we'll go into some of those guardrails and red lines that you've put in your report. I personally wanted to just unpack the recommendation for a moratorium on the use of high impact facial recognition technology. So for those folks that are watching this, that that sounds like a big deal. Explain what that is, and how that would actually change things at the moment.

Ed Santow 10:48

I should first say what we're not talking about. So many of the people listening in today will have smartphones that have facial recognition built in. But it's a relatively simple form of facial recognition. So is facial verification. So you use it, basically, instead of using a password or a pin to unlock your phone and it's really asking a basic question, which is, is the person whose face is in front of the phone, who they claim to be? So it's known as a one-to-one. We're much more concerned about a different type of facial recognition. That's one-to-many, in other words, picking the individual out of a big crowd of people and trying to identify them, especially when it's in a really high risk area, like policing. And what we've said is, there are two main problems with that. The first is, the technology isn't nearly as good as people think. I watched the same cop shows as everybody else. And it looks amazing on those cop shows. But the truth is that one-to-many facial recognition, on the whole is probably less accurate than anyone you know who is not blind. Just think about that. So it's much, much less accurate than normal people who can see. And so if you're using that technology to identify someone in a really significant situation, like, you know, you may prosecute them, you may detain them, all of those sorts of things. And then the technology has to be accurate. And you know, the errors are not evenly distributed. So it's much less accurate with people who have darker skin than people with physical disability. Even women, it's much less accurate than men. So that's one problem. The second problem, though, is even if the technology improves to some rate where it's perfect or as good as humans, and remember, it's nowhere near that now, I would still be really concerned. Because the more we use that sort of facial recognition in high risk areas, that the closer we get to surveillance state, I think we need to ask a really big question, which is whether that is what we want? And for me, the answer is no, we need to have really clear limits against misuse of facial recognition, but also overuse.

Peter Lewis 13:27

Lizzie, this has been something that Digital Rights Watch has been campaigning against for a while now. I guess you will fall into the camp that says we do not need this sort of technology at all. But where are those lines as far as you're concerned?

Lizzy O'Shea 13:41

Yeah, I think it's a really tricky question, because I think it's right to point out that the systems aren't good at this point. But the question begs then, well do we want them to get better, because then they can be deployed in very powerful ways. To give you an idea, I always come back to this example, because I think it's quite compelling one, Human Rights Watch did a report on this a little while ago, and they talked about how police, I think it was in New York, but in a major city, found footage of somebody committing an offence and in a bodega, it was a robbery and they ran the person's face through a facial recognition database that they use and no one matched and they thought the offender looked like Woody Harrelson. So they then ran a picture of Woody Harrelson through their facial recognition database, found some matches and actually arrested someone as a result. And I mean, it sounds ridiculous. But I do think this is what can happen in these contexts. So someone's been charged with an offence because they look like Woody Harrelson which is astonishing. But what it also shows I think, is that people who might appear in those databases already, because that's a database the FBI uses and others that's drawn from mug shots, you can see how people who are already maybe of interest to law enforcement and it ends up becoming a self-fulfilling prophecy where overly-policed populations get policed even further, because that person is already in the database. And that's my real worry as well. I don't think we want a mass database of everybody's faces to be able to carry out law enforcement purposes. And we don't also want to build it slowly over time on an as-needs basis. I think we do need to think seriously about stopping altogether the use of this technology in certain settings. And there are cities around the world that are doing this. Because part of the problem is policing is an obvious place in which it's, it feels like a real violation if it's misused. But also I don't think people like walking down the street and knowing perhaps that in their local community, there might be facial recognition technology or cameras that can be deployed to do that. I think that the physical spaces that we occupy, also need, there needs to be some freedom from surveillance and some freedom to go about your business without constantly being watched. And there are cities around the world that have imposed moratoriums for that reason. And there's also companies that are now no longer selling facial recognition technology to law enforcement. And so that kind of initiative, I think, needs to be encouraged. And that's why I think this moratorium is great. We think that we need to start setting some serious red lines where we shouldn't be using it at all, and to have conversations about perhaps where we want to have freedom from it, even if those risks aren't quite as pronounced as in, for example, fields like law enforcement.

Peter Lewis 16:16

Ed, what would need to stop if your moratorium idea is accepted? Are there, will this ruin Peter Dutton's next 12 months or is this more prospective technology that you're trying to kill in the bud now?

Ed Santow 16:32

If Peter Dutton thinks less about me than I do about him.... but it's certainly true that the government has a bill that has had a bit of a troubled passage. It's called the Identity Matching Services Bill. It's been before the Parliament now for years, and the powerful Intelligence and

Security Committee knocked it back. It's a really complex bill, but at its core, it's actually quite simple. What it would allow government and some corporations to do, is create this big virtual database of facial and other biometric information held by Federal, State and Territory Governments, and to use facial recognition as well as some other similar technology for a range of purposes. Some would be, you know, to try and catch criminals, some would just be general policing, whatever that means. Some of it will have no connection at all to policing, and well, potentially, so it certainly isn't ruled out by the bill as we last saw. In fact, I think there was an acknowledgement on the part of the government that it provides the legal framework that could ultimately enable mass surveillance. And so I think that that is something hugely worrying. I have professional views and I have personal views, that we don't want to go down that path; that the targeted, proportionate uses of biometric technology like facial recognition may be justified, but you need to consider them on a case by case basis. But the idea that government should basically have completely free rein to use this in any conceivable way, I think is not the way that Australians want to go. And it would be really hard to justify that under our human rights law obligations.

Peter Lewis 18:28

Going into another part of your report - you look at the automated decision making that government is really beta testing on the Australian public at the moment, and the first beta test was a pretty overwhelming failure with Robo debt. And your recommendation is that at the end of every government decision, there needs to be a human taking accountability. Again, what will you be disrupting if that is accepted?

Ed Santow 18:56

So really, all we're asking for is something really basic, which is that when governments and companies use AI automation to make decisions, that those decisions are fair, and that they're accountable. And really existing law already largely requires that. But let's, let's be honest, right, the existing law, we've had difficulty applying existing law, as AI has become so common. So I think the ask is absolutely a reasonable one. It shouldn't make a difference. I mean, you know, we've had these conversations before, where we've said, Look, you know, imagine a bank, you know, has a sign in its window, saying, we're not going to have customers who are women, we're not going to have customers who are of this particular ethnic background, or we're not gonna have customers who are Jewish (I'm Jewish, right?) We would all acknowledge that it's completely unlawful, completely discriminatory. But if an algorithm has that same effect, an algorithm prevents women or people of a particular ethnic background or Jews from getting a certain product being served. We suddenly question ourselves and we ask well, is that a a legal issue or is an ethical issue and I want to bang my shoe on the table and say, it's, it's a legal issue, it is completely unlawful right now, but we need to enforce those laws. And one of the problems with some forms of automated decision making in AI is that you're left in this sort of tricky zone, where you're, you get this unsettling decision. And you think, maybe the decision is unfair to me because of my race, or my age, or my gender. But I'm not quite sure. And if the decision is so opaque, you can never kind of get beyond "the computer says no" initial response, then that is a fundamental threat to the rule of law. It's something that our entire democracy is based on. So it's really important that you don't have completely opaque decisions by government or companies, so that you can make sure that those two basic principles

that everyone, you know, almost takes for granted in Australia - that the decisions are fair and accountable, that those principles can be upheld.

Peter Lewis 21:21

Lizzie, I was just gonna throw to you and say one of the weird things, I think, and one of the reasons Robo debt took a while for people to sort of push back on, was that it just seemed such a crazy notion that this would be used on people that were vulnerable.

Lizzy O'Shea 21:39

Yeah, I think rightly so. And there were points at which I think a different decision could have been made and it was kind of obvious that the technology wasn't working as intended, you know, that they, they knew before they deployed it, for example, that the number of notices would go up by tenfold. And yet they still carried on, they weren't concerned that that might mean that some people got an incorrectly generated debt. So I was interested to know what you thought Ed about some of these processes. I'm someone that's never worked in government so often, when I talk to people who have or who engage a lot with government, they claim that there are a lot of people in these departments who do want to do the right thing maybe or raising criticisms, they don't necessarily have the capacity to change the outcome. You do talk about human rights assessments and the need to do that, before you deploy a piece of technology. I mean, is that your experience of talking with people in government, that there is, I guess, a spectrum of people who are interested in making good decisions, in designing good systems, and there are perhaps people who are less interested in that? And these kinds of processes, which - you know, I'm a litigator as my day job, I like enforcing laws in courts - but I can see a system where having to go through those steps and do that assessment does then change the dynamic internally, in terms of the speed and the considerations that you might take into account when designing and then testing and ultimately deploying these kinds of automated decision making processes?

Ed Santow 23:06

It's a great question. And maybe I'm an inveterate optimist, but I generally think that people inside government and outside government don't, on the whole want to violate people's human rights, that that's, that's really unusual, right? What's really particular about this area, and it took me a while to get my head around this, is that on the whole, when government uses something like Artificial Intelligence, or even automation, they're not just developing the the decision making system themselves. In fact, usually, it's a partnership where they relied very heavily on people from the private sector companies in the private sector to develop the decision making system. And that creates, you know, that's outsourcing, right? There's a whole big argument about whether that's a good thing or a bad thing, but it certainly is just a reality. And it does create a real weak point here. Because if you don't have people in government, who are at least asking the right questions, in that process of procuring, you know, buying an AI system, and tweaking it so that it's appropriate for the particular government agency, then it seems almost inevitable that the system is not going to be tailored right for the particular circumstances. We've seen overseas, for example, when the US Justice System, you know, bought its infamous Compass system, that the company that was developing that system didn't understand that there was a correlation between where someone

lives in the US with the zip code, the postcode, and their race or their ethnic origin. And every Justice Department probably in the United States, and certainly here in Australia, would know that there's that correlation. But unless you, you're able to point that out, that if you're taking into account someone's postcode, you're effectively taking into account their rights, then you're not going to be able to develop a system that is safe. And it's, that's all I'm asking for, I'm not asking for an absolutely perfect system, I'm just asking for a system that is safe in the sense that it will uphold people's basic human rights. So that process of procurement and implementation involves these partnerships between government and the private sector. And that's why we need to really spell out the process so that can happen in a more sophisticated way. So these sorts of problems can be brought to the fore earlier and can be resolved. And ultimately, if they can't be resolved, in that design and development stage, they're not ready to be used.

Peter Lewis 25:56

Ed I imagine that the pushback from the industry is that "you don't actually get how innovation works". And innovation is all about throwing things out there seeing what sticks, learning from what doesn't, iterating - is this anti-technology this report? Is this a barrier for Australia to be a tech player?

Ed Santow 26:20

I mean, I think we sometimes fetishize a bit, that process that you've just described. In some areas, it really doesn't matter too much. If you're developing a computer game, you know, the stakes are pretty low. We're not too worried about what's going on there. By contrast, if you take this out of the tech area, if you're developing a new medicine, or if you're developing a new car, no one would say that that's okay. You just, you just put something out there, you know, it's not right. But you iterate, you see how many people die, you know, from drinking your potion or, you know, on the roads, because your, you know, your car acts as a, has an ejector seat, and then you fix it as you go. No one would say that's okay. In fact, it's the very definition, to my mind, of being like reckless. And so I think what we have to do is be able to differentiate between where a tech company is operating in the area that's genuinely low risk, and then, of course, the expectations on them should be low. But if they're operating in an area that's much higher risk, where if you make a wrong decision, people are really badly affected, then, of course, the expectation should be greater.

Peter Lewis 27:33

So you think that there will be tech companies that embrace this, because they can see it building a more robust way of doing their innovation?

Ed Santow 27:41

Absolutely, absolutely. It's really tricky for me, as you know, a Government Agency to name them, so I won't. But there are some really, really impressive tech companies that we've come across throughout this project that take their obligations seriously. There's sometimes a bit of a kind of a 'hare and the tortoise' type scenario here, where you can, you know, iterate, as you say, really

quickly and come up with beta products. So these sort of experimental products that you just sort of throw out into the community and see what happens. And you can make a bit of progress quite quickly on that. But if it turns out that you're causing huge harm, then you have a choice, right? It's a really difficult choice. Either you just ignore that and you just hope that no one's going to ping you on it, which is disgusting and terrible. Or you have to fix it after your product has already, or services, already kind of taken shape. And that's much harder, much more expensive, and tends to yield worse results. So we talk a lot about human rights by design. And there are a number of tech companies and others who take that process really seriously. So they speak with their community, the people that they're trying to serve, and they really, but they don't just speak with them I should say, they listen to them as well. So they ask questions, they try and involve them in the design and development of their products and services. And that may take a little bit more time at the beginning, but it tends to yield much better results. And they are the 'tortoise' that tends to win the race.

Peter Lewis 29:20

And that's where you get your Danish furniture.

Lizzie O'Shea 29:23

I mean, this is the thing. I think we don't want to be a world leader in scandals about poorly designed algorithms that punish vulnerable people. That's where we do not want to lead the world. And at the moment we've got Robodebt on a record and here's the opportunity to learn and execute properly because I think one of the problems in, you know, a vendor-driven marketplace, say, for providing these services to government is that if there isn't baseline expectations or regulations or requirements, those who do make an effort to go above and beyond what's required because they think it's the right thing to do, get penalised of sorts in the marketplace. So it's about lifting the whole market's attitude to these questions through that kind of expectation setting, whether it's regulatory or in, you know, human rights assessments as Ed advocates for. I think it's really critical to making a better marketplace for technology, as well. I mean, the other thing I wanted to ask you about Ed is whether you've encountered this issue where people who are making these systems then just say, "Oh, well, it's just too complex, we don't even understand how maybe it's arrived at this result." So therefore, kind of anticipating or trying to reduce these issues of harm that arise after the fact, it's just an impossible task, because that's the other thing I think sometimes industry says, it's..or sometimes government ministries, who have deployed, or government departments that have deployed this technology and perhaps, you know, have resulted in unintended consequences, that it was just too complex and their systems are too complex to understand. And I wonder if you've encountered that attitude and what you would say about it.

Ed Santow 30:58

There are about half a dozen technical questions that were incredibly difficult, that we felt we absolutely needed to answer, before we could finalize our report, and that was one of them. So they are, you know, kind of exotic forms of Artificial Intelligence, like big neural networks, right? Where the system is learning from itself, it's.... you have duelling algorithms, anyway I won't get stuck in the jargon. Let's just agree, they're super complicated. And it's true, that it's very difficult to understand

how those systems arrive at the results that they arrive at. So I spent kind of a couple of years making a bit of a pest of myself, asking literally the world's leading experts on this, is that something that can be solved? In other words, if you have really sophisticated AI, is it incapable of, you know, yielding explanation for the results that arise? Or is it a design question? Is it just more difficult to design the AI that way? And it seems that the better view among experts, is that it's a design problem. And thank God, right. So what that means is that, you know, the vast majority of AI out there, if you design it well can yield an explanation. Not a perfect explanation but humans don't give perfect explanations for their decisions either. But an explanation that's meaningful, and we'll get to where we need to get to. And so what that means is that companies have a bit of a choice: they can either, you know, go the extra mile, and design this system in a way that is capable of providing that sort of explanation. Or if they're not willing to, you know, invest in that, then maybe that they shouldn't be using that particular type of AI. And just before I move off that, the aviation industry is a really good example of this right. So you would think that in aviation, particularly designing new airplanes, that would be an area where they would be absolutely all over AI, and particularly that really super complex AI. But the truth is, they're not. And part of it is because of regulation that has served us incredibly well. So whenever there is an incident, and certainly, you know, whenever there's a tragic plane crash, there's an investigation and the airline has to be and the manufacturer has to be able to provide the authorities with the information that helps them get to the bottom of what went wrong. And if they can't do that, because they just sort of say "well a computer basically spit out this answer, and we just did what the computer said", that doesn't cut it. And that's a really good example of a high risk area where they just can't get the AI right to to use it in a safe way. And so they've made the decision at the moment, "well, we won't use it as extensively as we might". I think that's the right answer frankly.

Peter Lewis 34:09

Eb, do you want to go to the floor and see what's coming out of our fantastic audience?

Ebony Bennett 34:16

Yeah, we've got quite a few questions here. We've got, I think, more than 360 people on with us today. So thanks so much for joining us and for some of these great questions. The first one I was gonna ask is from Johanna Badcock. She says that facial recognition seems to have been rolled out in parts of Perth in Western Australia without community consent, and that the only way to get around it is more or less, to just avoid the areas where you know that's in place, and comments that it's an unreasonable restriction and, you know, access around movement. And I wonder if you could comment on kind of what the report sheds light on in terms of where these are already being implemented and how citizens like Johanna might respond where these types of systems have been rolled out.

Ed Santow 35:09

So what we say is that there just isn't enough data about how it's being used and there aren't clear enough rules. So Lizzie makes a really good point that we don't have a Human Rights Act or a Bill of Rights in Australia. If we did, that would put some guardrails around this, but of course, the UK is probably the country that uses the most facial recognition, including in policing, and they do have a

Human Rights Act. So it's not a panacea but it would make a difference. What essentially we're saying is, at the moment, there is insufficient protection against misuse of facial recognition, including by police in Australia. So that leaves our citizens very badly exposed. I'm not making any criticism of any police force; all I'm saying is, there are real risks with this technology, governments should say, at the very least, we're going to make sure that we get community buy-in before we use this. And also, as part of that process, we will put in place really, really strong protections. And I'm afraid that just doesn't exist right now.

Ebony Bennett 36:22

Yeah. Which brings me I think, to the next question that we've already kind of touched on, which is from Stephen Masters. He says, "Why doesn't Australia have a Bill of Rights? I wonder if I could just get a few reflections, maybe starting with you Lizzie?

Lizzie O'Shea 36:38

That's a good question. There is a campaign underway to change that. So I'd encourage people to, if they are interested in this topic, to look that up and join that campaign, because I think it is pretty critically important. And I agree with it, I don't think it is a cure-all for the problems that we face. But we do end up beginning at a standing start if we don't have it, because even the basic language of rights, how they work, the idea that they're not absolute, that they work, you know, that they're interdependent, that kind of nuance is lost when you don't have an instrument or a charter that you can point to. And it does also mean that government feels less obliged to comply in the ways that we were talking about before institutionally. What I would say just also, about this question and the previous question is that we, um, we at Digital Rights Watch talk a bit about, there is a UN Digital Cities for Digital Rights campaign. And there's a charter that talks about deploying technology in urban areas, and you know, why you need to have and commit to public participation and engagement to make sure people know what kind of technology is being used and they can have a say in it. And that, there's a document, a charter, that goes with that. And we think cities should sign up to it so if you are in Perth, you might want to talk to your City Council about that, because we've certainly written to lots of councils asking them to sign on. And in some ways I think it's an interesting phenomenon because, you know, we see this in places like the US as well, where there might not be federal laws for certain kinds of protections, but at a state or a very local level, that kind of campaign can start bubbling up and you can get, you know, well, for example, in the United States, they're not signatories to the Convention Against Discrimination Against Women, but cities have signed on to it. And I think we could do something similar in Australia, where local councils and even state governments who have already done this to some extent, start implementing these human rights, and that helps generate the campaign for a federal instrument as well.

Ebony Bennett 38:39

Ed, I might come to you on that question next. Obviously, there's always a lot of criticisms floating around for a Bill of Rights and I know the Human Rights Commission is based around a lot of anti-discrimination laws, in particular. How much of a difference would it be to protecting human rights if we did have a Bill of Rights in place?

Ed Santow 38:57

I think it'd be really significant. You know, at one level, it's philosophical. A country that doesn't have a Human Rights Act, if you've got a problem, particularly, if you've got a problem with the government, you essentially have to go to the government as a supplicant, as a beggar saying, "please, will you help me"? With a Human Rights Act, it changes the relationship between the government and the citizen and it says to the, you know, the person that has a problem with the government goes up to the government and says, "These are my rights; let's work out a way to make sure that those rights can be upheld". And that's a really fundamental philosophical shift. If we've been having this conversation sort of 20, 25 years ago, there would have been quite a few countries, so only in the Commonwealth, but there would have been quite a few countries that still didn't have a Human Rights Act or Bill of Rights. Now, I think we're pretty much the only liberal democracy in the world that doesn't have one. And yes, I think we all agree it wouldn't be a panacea, but it would make a significant practical difference. It's a little bit exhausting. The first proposal for a Bill of Rights in Australia was made in the 1890s. And every couple of decades, we have this debate again. I was very involved in the last process just over 10 years ago. I think, eventually, we will have one. But it's certainly taking a long time.

Ebony Bennett 40:35

So a good reflection, I guess, the last year with the pandemic has shown that huge changes in politics are possible at short notice, if enough people are willing. Um, the next question is from James Bannon, he says "one of the biggest issues with Artificial Intelligence is that models are often proprietary and opaque. And how can governments make platform vendors legally transparent and accountable without also stifling incentives to innovation?" Ed, I might come to you first on that one. And I guess also, after the experience of putting together this massive report, just a comment generally on how committed industry is, to being transparent, or is it kind of a secondary thought?

Ed Santow 41:18

Yeah, we'll have to wait and see a bit now. It's only sort of a day after the report was tabled in Parliament. So that's a bit too early to say. We certainly had, I think, very conscientious input and engagement throughout the process. The idea of, you know, opening up the recipe to the secret source is one that comes up all the time in Artificial Intelligence. And I don't want to suggest that, you know, they should just be open slather and companies shouldn't be able to protect their trade secrets at all. All I'm saying is that if something goes wrong, if there is a problem with.., if someone feels that they may have been treated unfairly, there must be a process where the decision making event can be properly scrutinized. And essentially, technology shouldn't be treated differently from every other area of human activity and told, "Well, you know, that that's completely secret but everything else is open to proper scrutiny". So if something went wrong, you know, in a hospital, and there was some, you know, very, very closely held secrets about the medicine that, you know, may have been incorrectly administered, a court could, of course, look at those secrets. And it would keep that information confidential. And we've got very effective processes to do just that. And so I just don't accept at face value, this idea that any scrutiny of how AI works, or how it is trained, will necessarily mean that those companies will be disincentivized, from engaging in that work.

Ebony Bennett 43:18

Pete, did you want to add anything to that?

Peter Lewis 43:21

Yeah, I think it's really interesting in terms of the broader discussion about algorithmic transparency, which is this whole assumption that things that are designed by our algorithms are somehow sacred and can't be subject to scrutiny when they have such an impact on the way - whether it's our information systems, whether it's a government service delivery - is rolled out, and I'm not sure Ed whether your idea is that "the algorithm's there" and under challenge, it's discoverable or, you know, if you're building a new car, as we said before, you've got to show the plan. So you need to get that through safety checkers. So it might go into this role of an AI Safety Commissioner, that you've got in your report. Is that sort of your thinking of where the scrutiny meets the algorithm?

Ed Santow 44:13

Possibly. The interesting observation, which I've mentioned to you once, before Peter, is I had this realization, and it took me way too long to realize this, but I'm going to make this confession in front of all of you. I realized that, on the whole, AI is talked about in these magical terms, but it doesn't really allow us to do many things that are wholly new. Right? It really doesn't. It allows us to do things that we've always done but in new ways. In other words, we've always made decisions in banks about who's going to be a good customer, we've always made insurance decisions. We've always, you know, made decisions in government about who owes what money to the government. But AI creates this new way of reaching those decisions. And so what is critically important is that, however those decisions are made - whether they're made using the most old fashioned abacus, or the most sophisticated AI - that they are open to scrutiny in, you know, essentially the usual way. So it's really on the whole the existing regulators that need to step up and say, "Look, if you're using AI to do this thing that we're responsible for, you've just got to comply with the laws that are there to protect the system. It's that simple and sometimes I think we can tie ourselves in knots by saying, "Oh, well, there's been AI sort of sprinkled like holy water over this process". And we've suddenly say, "Oh, well, then okay, in that case, we've got to step right back". Why? If it's a robust system, and if it delivers good, safe results, then we're not asking for more, we're not saying that there should be greater scrutiny of it, we're just saying that the normal scrutiny should apply. And that means sometimes, before you put the system out, you have to, you know, show your workings to the regulator. But certainly if something goes wrong, or someone claims that someone or something has gone wrong, then you have to show your workings for it.

Lizzie O'Shea 46:19

But also, if the best protection against having to disclose it in that context, whether it's in a court proceeding or for some other, in some other avenues, to not have those mistakes happen, that's the other way in which...

Peter Lewis 46:31

Don't crash the car.

Lizzie O'Shea 46:33

Yeah, or don't, yeah, don't make a car that explodes or, you know, that's the kind of thing. It's not just also a stick, I suppose, it's a carrot too, to say, well, if you are designing these things, and you've done it well, you will not have harm... you're protecting yourself as much as you can against harmful outcomes. And here's a methodology for doing that, which is employing a Human Rights lens to think about who's using these systems, think about people as citizens with rights and dignity that ought to be respected rather than just as guinea pigs to be experimented on. And if you keep applying that process throughout the stages of design, you'll decrease the chances that you will have a harmful outcome and have to expose that reasoning that you've done. And if you have to, that it will be well documented. So I think that it works, that kind of regulatory approach for that expectation is good. And we don't let then companies get away with saying, "Oh, well, this is just very complicated. And, you know, we could never have anticipated this happening". I think the days of that kind of defence are well over.

Ebony Bennett 47:39

That takes me to our next question from Thomas Denise, who says, "Should manufacturers of AI be required to put their systems through a series of trials, like pharmaceutical companies before sending out a new drug?" Does anyone want to comment on that?

Ed Santow 47:56

And the short answer is, it depends what they're using it for. On the whole, if you're, you're not really manufacturing AI, you're manufacturing a product or service or a facility that uses AI. So, as I say, if you're doing something that's really low risk, like, you know, creating a new computer game or something like that, then the expectations, the regulatory expectations, are much lower. But if you're doing something that's high risk, if you're using it to create a new medical drug or something, you know, a credit decision making system for bank loans or for Centrelink or for the NDIA, as is currently proposed, then you need to be much more rigorous in the testing process. And you need to be much more open to scrutiny if people claim that something goes wrong.

Lizzie O'Shea 48:52

I mean, that's what they're doing in Europe, right? Ed, I mean, they're setting up a distinction between high risk and low risk technology, and then setting the standard based on that, which I think is interesting, because we appear to be on the same page. So that's good.

Ed Santow 49:06

Exactly, exactly. Because on the whole, you don't really regulate technology. Sometimes you do, sometimes you absolutely need to. We've always regulated nuclear technology. But as a general principle, we tend to regulate how technology is used because that's where it is most important. And that's where you can really determine what the risks are.

Ebony Bennett 49:31

This might be a quick one for all of the panel to reflect on. It's from John Flacker. He says "How have we done in terms of educating policy and decision makers on both the capabilities and the risks of applied AI based solutions? And if we haven't, how do we address this?" I might start with you Ed.

Ed Santow 49:52

I think we've still got a bit of a way to go. You know, in a sense, you can - I'll try to be really quick - You need to sort of create three divisions. So people who are making really big policy decisions, or, you know, company, chief executives or equivalent in government, they really need to understand those big strategic risks and opportunities and where AI is reliable, where it's not. One layer down, the middle managers, the people who are responsible for actually implementing these systems, they need a lot more information. And then the rest of us just need the basic information to protect ourselves. And so as AI is increasingly going to be part of our world, we need to provide that sort of level of education and training at each of those three levels more effectively.

Ebony Bennett 50:42

Lizzie, did you want to comment on that?

Lizzie O'Shea 50:44

Yeah, I think it's a really interesting question. Often when we talk about, when I talk about technology and human rights and these kinds of topics, people lament that our politicians don't really seem to understand technology and that's how we get these bad outcomes. And I understand where that sentiment comes from but I think we should be a little cautious about that. I think there are politicians that know exactly what they're doing, or have deliberately outsourced the decision making to people who do. And so I do think we need to hold our politicians to account more on bad tech policy outcomes. Having said that, I appreciate that these questions are kind of complex and difficult to contend with it. Ed just talked about industry and I think it's really interesting. I saw someone writing about the drive to recruit people in the tech sector, who have expertise in ethics. And this person is a tech industry person, was mapping out all the ways in which you need somebody to be able to have those skills at basically every level of the process of creating products, and just have complexities that actually, when we talk about, you know, making tech ethical, or 'tech for good', that is not just about how you design the ones and zeros and that kind of thing. It's also thinking about your audience, you need to have some context and history, some understanding of the politics and the environment that you're engaging with. You also need to have people who understand security and how important that is to the people who you're designing for. And you do realize just how how multi dimensional it is, and how it needs to happen at every single layer. And so

I do think there are companies that are working towards that. But I think there are probably it's fair to say, companies that have preferred to avoid and shirk that responsibility, because they perhaps grasp how complex it is and can see the potential for it to disrupt the way that they've been working. So I think in both instances, we need to keep the pressure on, because the age of kind of willful blindness to these problems, I do think is, has come to an end. And it's our job then to hold them to account.

Ebony Bennett 52:40

And back to you Pete.

Peter Lewis 52:42

Yeah, look I think two things. One is that tech policy across government is really fragmented. So Ed's reporting into the Attorney General, Communication's minister's got a lot of the interaction with the platforms, all the ACCC stuff goes through the Treasurer's, there's a Minister for Digital Transformation, I believe. There's also Cybersecurity, another department, and E-safety, another department. So the government hasn't got a 360 degree so there's no sort of overriding push to make that happen. And hopefully, this report can address that to some extent. And I think the other thing from today's discussion is how we frame this up not as a constraint on technology, 'oh, my God, we've got to be ethical and fair', but actually how we make this the great secret source of Australian technology, because it's ethical and fair, it's going to be better. And it's the sort of technology we can be exporting to the region and the world as something that's an alternate to the hyper-capitalist Americans and the state surveillance of the Chinese. So it's an optimistic story around this and I do think this has been a contribution to it.

Ebony Bennett 53:56

Excellent. Well, Pete, did you have anything else that you wanted to kind of address before we wrap up?

Peter Lewis 54:01

Yeah thanks Eb. I just got the last laugh, particularly from Ed but also Lizzie. So Human Rights Commission reports have a great tradition of ending up in the circular filing cabinet underneath the minister's desk. What do we need to do as this audience and beyond to make sure that this one is the one that leaves and actually moves to the next level?

Ed Santow 54:27

And it's certainly true that our reports tend not to be embraced immediately. But sometimes five to 10, often 5, 10 years down the track they're just part of policy. The issue for us is we don't have five or 10 years. We really are at a critical point now, where AI is on this exponential rise, and we need to get to grips with the kind of society we want to be. We want to live in a society where new

technology like AI gives us what we want and need, not what we fear. And so it can't just be us saying it. We've relied very heavily on this two stage consultation process here and overseas, listening to experts, like you Peter, and like you Lizzie, to help us frame these recommendations. But now, this is, not just a gift for the government, it's a gift for, hopefully, the Australian people so that people can say, well, these are the things we really, really insist on. And if I may draw an analogy, you know, we saw that last year, my colleague Kate Jenkins, the Sex Discrimination Commissioner, she published a very, very important report "Respect at Work" on sexual harassment, on a very different subject, but a very important report. And really what it took was a strong, I guess, wave of community concern, saying, "We don't accept what's happening at the moment". And as it happens, the government has some really strong advice from the Commission about how to solve that problem, to start to lever change.

Peter Lewis 56:14

Yeah. We did do a bit of pre-emptive polling, not that we saw the report in advance, but we did find that there is widespread support for some of your recommendations. And how do we tap that general support for regulation at the moment, Lizzie, to make sure that this isn't just another report that sits there waiting for the next catastrophe?

Lizzy O'Shea 56:33

Yeah, I think that's a really interesting point to make, that lots and lots of people in society, I think, would think that what is in that report that Ed's done, is all very sensible, needs to be implemented, if there wasn't an expectation that it was already happening. And so it's translating that into genuine policy change and, you know, at Digital Rights Watch, we like to obviously run campaigns, we've got a campaign against facial recognition technology, we've a campaign for Digital Rights Cities, we do other kinds of campaigning work. We also try to lift people's literacy and how technology is used in their lives and how they can also resist some of those uses, for example. I think it's also about building a campaign that's as diverse as possible and I suppose that's what I wanted to offer - that we certainly think at Digital Rights Watch, that digital rights are not just something for technologists to be worried about, or people who've got experience in dealing with tech. That actually it's something that people in lots of different campaigns, whether it's a campaign, you know, to build a civic in urban spaces, or to campaign against violence against women, to advocate for the rights of people with disability, which is also a significant section of Ed's report that we didn't get to discuss today. But all these organizations that work in these different spaces have an interest in making sure that we get these technology decisions right. And we want to involve them in it. So that's my kind of approach to these problems - build a movement that's as diverse as possible, figure out ways to communicate to people and to listen to them as to why technology and technology policy issues are so critical in these other fields and then, from that diverse movement, we can have a strong and compelling message to the government that they need to act on this. And also that industry needs to act as well and lead, so that we can avoid harm, we can have technology that's about human flourishing, not just oppression or surveillance, and we end up having a 21st century that's much more promising than it might seem at this particular point in time.

Ebony Bennett 58:40

Thank you. And I can see a good response in the Chat here that the Jenkins report that Ed mentioned, was prompted by a long list of victims of sexual harassment and other problems, and that we shouldn't wait for victims of Artificial Intelligence and these other automated decision making processes to implement the recommendations from Ed's report. So we will wrap it up there. Thank you so much, to all three of you for joining us today and for such a fascinating discussion. So that's thanks again to Ed Santow, Pete Lewis and Lizzie O'Shea. We really appreciate your time today. And thanks, everyone, for joining us. So some really cracking questions in there. I'm sorry that we didn't get to all of them. But if you are interested in these tech discussions and these tech talks, they have been at Australia At Home fortnightly for quite some time now. And the Australian Institute Centre for Responsible Technology is hoping to partner with Sydney Policy Lab at the University of Sydney in future. So maybe drop us a line to webinars@australiainstitute.org.au if you'd be keen to hear about those in the future, and make sure that you're subscribed to our podcast "Follow the money". You can find that on iTunes or wherever you normally listen to podcasts. Thanks so much for a fabulous discussion today, everyone, and hopefully we'll see you again soon. Thanks.

Transcribed by https://otter.ai