Transcript

Title: Technology x Society Forum: Implementing Australia's AI Ethics Principles
Creator: Swinburne University of Technology
Year: 2023

Audio/video for this transcript available from: http://commons.swinburne.edu.au

[MUSIC PLAYING]

PAUL LAVEY: OK. We might make a start. Good afternoon, everyone. I'm Paul Lavey, the research events coordinator for Swinburne Research. Thanks for joining us for today's webinar, which is Implementing Australia's AI Ethics Principles which, is presented by both La Trobe University and Swinburne Social Innovation Research Institute.

You'll shortly hear from associate Professor Sam Wilson who is a program leader for Public Interest Technology at the Social Innovation Research Institute. But before I hand over to Sam, I would like to acknowledge the traditional owners of this land.

We acknowledge that we are hosting this webinar from the lands of the Wurundjeri People of the Kulin Nation. We also acknowledge the traditional custodians of the various lands on which you all work today and the Aboriginal and Torres Strait Islander people participating in this webinar. We pay our respects to Elders past, present, and emerging and celebrate the diversity of Aboriginal peoples and their ongoing cultures and connections to the lands and waters.

So at the end of the session, we invite people to-- comments and questions. So just please look at the Q&A box on your screen. Think of any question at that point. Just type it in there, and we'll respond to that after the presentations.

And just letting you know that we are recording the session. And this will be published on our website, and we'll distribute it to everyone at the end of the session. And at this point, I'll hand over to you, Sam, who will start today's presentation. Sam, over to you.

SAM WILSON: Thank you very much, Paul. And good afternoon, everybody. So my name is Sam Wilson. I'm an associate professor of leadership in the Swinburne Business School. And with Lawrie and our colleague, Diane Sivasubramaniam. We are the conveners of the Technology by Society Forum. The core focus of the Tech by Society Forum is public interest technology.

And for those of you who don't know about these types of technologies, public interest technologies serve to address social challenges and advance the public interest or the common good. Public interest technologies put people and society at the centre of our technological choices and strive to ensure that the benefits of technology are widely shared.

However, it's important to note that public interest technology isn't one thing. Moreover, it's not just about technology experts. Rather, it includes all those who are responsible for adopting and implementing disruptive technology, particularly in the public sphere.

The Technology by Society Forum, which is an initiative of Swinburne and La Trobe universities, provides a space for engagement between STEM and HAAS academics with an interest in the design, development, and application of public interest technology. Our aim is to stimulate collaborations between STEM and HAAS academics.

By examining the shared and distinct ways that STEM and has researchers think about technology and society, we seek to discover new insights into how to address complex social challenges and how best to apply that knowledge to developing research partnerships with industry and communities.

It's my pleasure now to hand you on to Professor Lawrie Zion from La Trobe University, who will introduce today's speaker, Rita Arrigo.

LAWRIE ZION: Thanks so much, Sam. And thanks, Paul, to you, too. And welcome, everyone. It's great to have so many people along today. I'm Lawrie Zion. I'm the associate dean of research and industry engagement for the School of Humanities and Social Sciences at La Trobe.

Rita, welcome to you. And thanks so much for making the time to talk to us, to this cohort of people from La Trobe, and Swinburne, and elsewhere. Really, your current role is a new one, strategist engagement manager for CSIRO's newly launched National AI Centre. And the centre was only just launched in March this year.

Rita, I know you've had a really diverse career within the tech sector, which you will tell us a bit more about in a moment. But yeah, just before we go to Rita, I want to reiterate what Paul said. And that is if you've got any questions or comments, pop them into the Q&A at the bottom of the screen, the Q&A tag, and we'll get to those in the second half of the session today.

But Rita, I'm going to hand it over to you now.

RITA ARRIGO: Thank you so much, and thank you for the wonderful introduction. And it's such an honour to be here today. It's really exciting to have these multidisciplinary initiatives like the Technologies and Society Forum, and I'm really honoured to be sharing with you today.

And I wanted to talk a little bit about the work that we're currently doing at the National AI Centre. So I'm gonna be talking about implementing Australia's AI ethics principles. Not an easy task, but I'm going to take you on that journey of what we're trying to achieve through that space.

But just to tell you a little bit about what I'm going to be talking about today, I'm going to share a bit about my journey, where I've been, all about the National AI Centre, the Responsible AI and Responsible AI Network which was launched in March, a bit about the AI industry in Australia, and how you can get involved to be part of the Responsible AI Network.

So a bit about my journey. I've been at this for a while. I have my first job in PC support, where I helped secretaries move from typewriters to PCs. And back then, there was a lot of challenges. Am I going to lose my job? But I'm telling you none of those secretaries wanted to ever see a typewriter after they saw a PC. [LAUGHS]

And throughout my career, I've been involved in a lot of emerging technologies. Like, I was involved in the first internet cafe in Australia. I had a radio show on RRR that's still going called *Byte Into IT*, which was about reducing techno fear.

But part of the internet cafe introduced me to the internet, and I became completely obsessed with what was possible. And I ended up working for Telstra during the big pond era. And from that, I worked for digital agencies. And that really got me into the whole digital transformation space.

So I ended up getting a role at Microsoft as a chief digital advisor, where I would do a lot of work with transforming a lot of our accounts into the new digital world. And it was at that time that I learnt about AI. And I started hanging around with the researchers and looking at some of the APIs, and learning about computer vision, and learning about machine learning.

And it really opened my mind to what was possible. And that led me to join the CSIRO National AI Centre, where one of the main initiatives has been this Responsible AI Network.

And I also want to say that I was lucky enough to get a bit of an AI education as part of judging the Women in AI Awards. So it's finally an area where women have been very-- not as well-represented as men. You get a chance to learn about all the amazing things that women are doing in AI.

The other thing that I learnt at Microsoft was mixed reality and holographic computing. And I still think that that has a lot of capability in the metaverse and AI. It's really going to change the way look at the world. So that's a little bit about me and where I've been. So let's talk a bit about the National AI Centre.

So the National AI Centre was established by the Australian government, and it's coordinated by CSIRO. And we have our foundation partners, Google and CEDA. And we're here to build Australia's responsible AI advantage. We need to enable responsible AI across industry and grow a responsible AI industry in Australia.

Super exciting, because I think it's an area where as Australians, we're really into that fair go and being able to make sure that the way we implement technology is fair and just. And I'm going to tell you more about the kind of areas that we've been getting into.

So the three areas that I'm working on is around getting started, like, helping the community understand what AI is. How do you get started? How do you get involved? What are some basic understandings in that space?

Getting connected. So we have an ecosystem discoverability portal, where you can actually find out about some of the AI startups that are in Australia. And I'm going to-- and organisations that are also beyond startups. They're like real organisations and bigger than startups. They're delivering AI

solutions. And I'm going to tell you a little bit about them as well and then the Responsible AI Network, where we want to uplift practise.

And at any stage, if you want to go back and visit any of the content that I'm talking about, can visit us at csiro.au/naic. And for those who know us well, we call it NAIC. So that's the National AI Centre.

So where are we today? There's this-- and I'm sure everyone feels it. I know I was talking to the panellists earlier today. There's this AI wave that's accelerating and happening, but we have these concerns around trust, privacy, security, and data quality. And we also see the potential of AI to alleviate many challenges from health to environment to societal innovation. But with all this complexity of everything generating data, we do have this moment in time to engage in responsible AI.

But we also, as I was explaining, we have this all these complexities around the world. And CSIRO has megatrends around adapting to climate change, leaner, cleaner greener, the escalating health imperative. And we see the possibility that I could be our co-pilot in this complex world and help us achieve some of these big challenges that we have to address across the decades.

So the moment in time to have responsible and inclusive AI is here, and that's what we're working on. And we're looking for lots of partnerships to get involved in this space.

And I do want to reiterate that AI does have the potential to deliver a lot of benefits, and it's not just me saying it. It's been predicted by a range of specialists. And it's also been very obviously noticed by industry. So it helps us to create better customer experiences, no more being on hold and not knowing what you need.

It helps improve decision-making. It can innovate our products and services to be more innovative and achieve cost savings. And most importantly, help us with productivity, which is something that we're always looking to do and improve.

But the other thing is that at the same time as all these benefits, organisations are not cultivating trustworthy AI. And this global index that was delivered by IBM shows that we have a lot of challenges in this space, around 74% not reducing that unintended bias, 68% not tracking performance variations and model drift, 61% not making sure that they explain AI-powered decisions. And there's a very topical example that's currently happening at the moment, which I'm sure everyone's heard about with robodebt and those kind of things.

Not developing ethical AI policies, not safeguarding about data privacy through the entire lifecycle, not monitoring AI across the clouds, not guarding against adversary threats and potential exclusions that keep system health and not tracking the data governance.

So there's a real need to think about, how do we do this? And if we're going to take advantage of AI, how do we make it trustworthy?

So in 2019, Australia was one of the leaders in developing our AI ethics principles. And so responsible AI is the development of intelligent systems, according to fundamental human principles and values.

And so Australia's AI ethics principles include the concept of human societal, environmental well-being, human-centred values, fairness, privacy, protection, and security, reliability and safety, transparency and explainability, contestability and accountability. And these are all very important principles. And I suppose what we're trying to do is find ways to turn these principles into practise.

So why do we need responsible AI right now? I think we have not had a day where we don't see the press around. Facial recognition is only accurate if you're a white guy. Fines for privacy breaches are soaring.

Even GPT, which was so popular-- there are people who are worried about copyrights. Stable diffusion is having a whole lot of challenges around artists' rights.

Microsoft, my old company, actually pulled out their facial recognition off the market. And you have to apply to use it and make sure that it's aligned to their ethical principles because of the way that people were using it, in a way that was not aligned to their responsible AI principles.

So there's always something happening in this space. So we really do need a way of being able to understand this area. And so the responsible AI network was designed to bring all these multidisciplinary parties in and to start thinking about, how do we deliver responsible AI and include some core pillars around law.

And so even though we have existing laws, there may be interpretations that may be required and thinking. Standards. So standards is really, really big. So we're working with Standards Australia. And I'm going to talk a little bit about the AI management standards that are coming out.

Design and how you design the whole experience. Leadership. And I'm so excited. That's why leadership is one of the most vital things that we need in this space. It's no longer OK to say, just talk to my tech people. It's about being able to understand what you're doing with your AI.

Principles, technology, and governance. And these are all areas that we're investing in to uplift responsible AI around these pillars and working with some core knowledge partners in that space.

We're also working on this Australian AI industry capability. And if you visit our ecosystem directory, you will see over 170 Australian companies that are working in this space, ranging from companies like Gaia, which won the Nasser award for being able to use AI to grow microgreens in the smallest space and are now in the Docklands, looking at ways that they can commercialise their product into different spaces through to people like MISH, who have developed a chatbot to help you with the escalating challenges around mental health. And I'll talk more deeply about some of these partners.

One of the key partners is CSIRO Data61, where they've been doing work on the Sydney Harbour Bridge by putting over 2,400 sensors on the bridge and combining it with machine learning and predictive analytics so they can tell engineers when and where a fracture will occur. They also have a range of robots continually monitoring and looking after that bridge.

Another organisation that's based here in Melbourne is Sapia. And Sapia has combined psychology, thinking, and AI to deliver a way of dramatically improving the way you recruit staff. And they've had amazing results, where they're finding that they are more likely to hire women because of the bias

that's being removed from the way that they do their questioning and answering through the chatbot.

Another organisation is HIVE, and this came out of WA, where health is developed-- delivered in a virtual environment, where they have a patient risk level alert system linked to a remote command centre using AI technology that monitors patients' vital signs around the clock. And this has dramatically improved productivity for the health sector and something that also improved the way health is being delivered.

And finally, one that's one of my favourites is Apate. And what they're doing is they're actually using conversational AI to combat scammers so that they'll actually put the AI against the phone scammers to make the conversations longer, to make it harder. And I just think this is amazing, to find AI to solve some of these challenges.

And as well, Camille Goldstone-Henry, who was one of the winners of the Women at AI Awards last year, a trailblazer and an Indigenous woman. She has been really focused on preserving biodiversity. And through her technology, she's been able to use the power of AI to increase conversation and impact and mitigate biodiversity loss.

So there's some key examples, but let's talk a little bit more about what we're doing with RAIN, so the Responsible AI Network, another acronym. RAIN. We have been delivering a range of webinars, and they've been super popular. And you can go back and watch them.

So we talk about the introduction to standards of AI. We've also talked about responsible AI engineering, which is a really big initiative that's happening within CSIRO Data61, where we've launched a patent catalogue and ways that you can take your engineering to the next level and make sure you're using responsible AI in that space.

We also talk about things like business analytics and the way you collect data around how your solution will work and how you can remove the bias in that area. And we have a range of other ones that we're also delivered last week, which is around the responsible operation of generative AI, which was super exciting. So much to learn about generative AI. And it's really broken down in terms of, what are the risks? And how can you avoid some of those risks?

And then we also had AI concepts, terminology, and the AI lifecycle which really went a lot more deeper into the AI management standard.

And we've also recently launched this paper, which I think if anyone is looking at ways that you can implement Australia's AI ethics principles, it's a super exciting resource. And we're not the only ones that think so as well. We've had it widely reviewed by The Australian, to the radio, the different media outlets that have really picked that up and seen how important it is to Australia.

And one of my favourite things about this paper is the infographic at the end, where it shows you all those principles I talked about, like human and societal, environmental well-being, reliability and safety. How do you do it?

And so in reliability and safety, you have to learn how to curate your data sets. You have to conduct pilot studies to monitor and evaluate continuously, and it gives you a whole lot of hints around that.

But it's also broken down into different personas, so if you're the system owner, or a senior director, or a development team. And I'm so glad that the senior director is being brought into this. Because these are the areas that directors can make a really big difference to the way we implement AI.

For example, in the human-centred values, directors can have a massive influence in incorporating diversity. In contestability, they can have an influence in supporting impacted individuals. So it's a great paper if you want to get your head around some of the ways that you can move from these principles to practise. And I'm hoping we get some questions about it, too.

So consumers are changing their awareness on what is fair. And you will notice that there's a lot of documentation around this. And consumers-- you will have noticed people, the consumer debate around putting surveillance cameras into Bunnings and how that didn't really need to happen and a lot of the other ways that we can improve that side of things.

And so there's a lot of consumer focus around privacy and data use, work, health, and safety, duty and care and negligence and consumer protection and anti-discrimination. And this is a great paper you can also have a look at that was delivered by the Human Technology Institute.

So not only is human's ability to grapple with some of the areas that AI is impacting, AI governance has also changed. And there's so much happening in that space with big data and bias in AI systems and the governance of IT and more and more standards coming out in that space.

And we're seeing a governance really being similar to the way ESG has matured with avoiding the AI governance by connecting AI into organisational governance processes. We are, once again, calling upon leaders to look at ways to put that into their ESG reporting.

And there's a great document that's also available around AI governance assets that you can look at and some of the top five supports required to implement effective AI governance from creating an AI strategy.

So it's not just about eliminating the risks. It's about thinking about how you can make your AI strategy actually provide benefits to your customers and people. Greater strategic expertise across the use of AI across the organisation. Greater board awareness of the potential risks documentations of policies and examples of effective AI governance and peer organisation. So a lot of work being done in that space.

And this has actually made the penny drop for me when I first started learning about, what are the things coming in responsible AI? And this is the AI model lifecycle that Standards Australia are using to implement the international standard for AI management.

And it shows all the steps in AI, from when you come up with the idea through to how you evaluate that idea, how you put the model in place, where do you get the data from. How is the data stored? How do you design and develop it? How do you deploy it? What's the user interface like for your users, and how do you verify and validate that?

And it's been super exciting to be part of what Standards Australia is doing because one day, you'll also be able to have a stamp from JTC that you're actually fulfilling all the requirements around this standard.

And also, I wanted to mention, when I was talking about leadership and how important that side of it is-- this was the study that was recently done by Fifth Quadrant and the Gradient Institute. And we found that when the CEO is involved, AI solutions are more responsible and they're better. So I think the challenges around educating yourself and being part of this area really do deliver results, from what this study showed.

I've got a few more minutes to go. And I know we've got a lot of questions coming. Am I right, Lawrie? Yeah.

LAWRIE ZION: I can't tell from where I'm sitting but if people have questions or comments start typing them into the Q&A panel at the bottom of your screen, I think.

RITA ARRIGO: Because I did want to--

PAUL LAVEY: No, we don't. I was just going to say, really, we don't have any questions yet. I think people are just taking it all in. And yeah, pop your questions in the Q&A when you're ready.

RITA ARRIGO: OK. I will keep going.

The other thing that has changed a lot of the way leadership works is generative AI, and we talked about that briefly. 30% of professionals have tried ChatGPT at work. 68% are using AI tools without informing their organisation. Globally, they saw that the company is considering ways to use generative AI. So there's a lot of work happening in this space, and it really has changed the way we look at governing AI.

Since we've got a little bit more time, I'll spend a bit of time talking about generative AI, if you like. Or I could just let you know--

LAWRIE ZION: No, Rita. Please. We'd love to hear any thoughts you have on the journey of generative AI. So yeah.

RITA ARRIGO: OK. All right. I'm getting some claps. This is cool.

So just so you've got some of the links, csiro.au/naic, NAIC. Follow us on LinkedIn for the latest information in that space.

So the fair go concept around responsible AI, we really see this as real. We're an early mover in the launching of Australian AI ethics principles. We've got a national, nationwide responsible AI researcher investment at Data61. We have the world's first e-safety commissioner in Australia. And so we really see this as something that we really can take to market.

We also have Responsible AI Month, Australia's AI Month coming up in November through December. And you'll see a lot of activity in a lot of ways you can get involved. So hopefully Technology Forum will get involved as well.

So a little bit about OpenAI. So Open AI was really the first with their launch of DALL-E and taking their ability to generate images and training neural networks to have image recognition and then be able to mash them up and then, obviously, with ChatGPT.

And I think there's not very, very few people that haven't tried it. But there are people that are using it more often than others. And this is me just asking it, how do I plan an event? And be able to get this real amazing detail out of ChatGPT.

But it's not just OpenAI that's using ChatGPT. We now have Google, who've launched Bard. And Bard is slightly different, because it actually doesn't just have a curated data set. It actually searches the whole of the Google ecosystem. And they've also launched the ability to build generative AI apps, which is a super exciting way of being able to have this app and be able to bring in a lot of these features.

Meta AI is also doing a lot of work in the generative AI space. And Adobe Firefly-- you could probably see it my background here-- generates, which is a new generative way for Adobe to be able to bring this stuff to market as well. And no longer having to have all those amazing Adobe skills, you can actually start to use some of the generative AI in that space.

And it's not just the big companies. So companies like Snapchat have also taken the ChatGPT API and implemented it into their app. And so people like that might just be using Snapchat, like all the young people are. I know my daughter is. All the students can now just talk to Snapchat and speak to my AI buddies, actually talking directly to the ChatGPT API. And pretty exciting for being able to get personalised lessons, et cetera.

And Duolingo is also using it to make it much more real, in terms of how you can go in and pretend you're at a cafe. And you get that generative AI response through that process. So that was just a little bit of who's using generative AI.

So let me just go back a bit. And Lawrie, I know you did have some questions for me. So why don't we get started?

LAWRIE ZION: Yeah. Thanks, Rita. Just such a fantastic presentation with so much to think about.

I think with all these changes happening really quickly, I'm wondering, someone, Rita, who's worked in the tech sector for a long time, firstly, does this feel like what's happening is different in kind from other rapid advances in tech? And much as I don't like asking double-barrelled questions, I'll follow that-- [LAUGHS] you can combine these answers in any way you want-- with a question about what worries you the most for all the good things AI can enable. What worries you the most about what's happening now?

RITA ARRIGO: Yeah. So I have been around for a while, [LAUGHS] doing this stuff, particularly in the emerging technology. And I think digital transformation was one of those areas that was super exciting, because you could finally see the web and everybody was using it. But even then, you'd walk into--

I had to work really hard to engage the business in this space. And it was something that was a real skill, because you had to take away the technical garble and turn it into a more business language. But very, very often, you'd get from business people, don't talk to me about this. Go talk to the technical people.

And I think sometimes that that is something that is not something that you could do with AI. I think AI has such a broader impact than the digital transformation, which was more about digitising paperwork and processes and as well, creating those self-service moments and those kind of things.

But AI is actually going to generate a lot more knowledge and interpretation and be able to see the world and have that computer vision element to it as well. And you start to think about autonomous systems and robotics and mixed reality. And it could also result in upgrading a lot of technology.

So to me, one of my greatest fears is that business people aren't going to embrace and feel that. It's too hard to understand, et cetera. So I think that's an area of super innovation, for leaders to be able to have their view of what's going on about it. So I think that's why the responsible AI network is so important, to be able to put it into plain English but also have those avenues for multidisciplinary conversations.

LAWRIE ZION: Yeah, Rita. I think I honestly can't remember if it's in the CSIRO publication or another one. But ahead of today, I was reading that on average, from a survey, 82% of respondents believe they're taking a best practise approach to AI. But on closer inspection, only 24% are taking deliberate actions to ensure that their own AI systems are developing responsibly.

And we're talking a lot about standards and identifying best practises, but how quickly is this gap between believing you're doing the right thing and not really being on top of it going to come back to bite us? And when I was talking to Sam before, he said, yeah. How many horses have bolted so far?

RITA ARRIGO: Well, sometimes I wonder whether we are in a Napster moment with ChatGPT, right? Because so many challenges around copyright and the way the data is being created and the hallucination and the security and all those aspects. And is it just an experiment that we're learning from?

That's because, I think, watching even at CSIRO, we don't put confidential data into ChatGPT. And I wonder how many people are in that position at the moment. So that's just an example.

But also, all the areas around how you collect data, how you store data. How do you, Diane-- how do you look at privacy? How do you do all that side of things? It's still very challenging.

And we don't want to leave it to a situation where rather than trying to address it now, you end up being in a legal scenario. And then it's up to the judges and the lawyers to interpret what happened and interpret all that side of things. You'd rather it happen before something terrible happens and those unintended consequences happen.

So I do think that there is a lot of work to be done in that space. And making time for it and investing in that is super important, because we also don't want to get behind in our productivity. So we want to be able to do both.

And I think that that's also part of the ability to be able to recognise organisations that might be already doing responsible AI is also really important, because a lot of organisations might buy off-the-shelf AI solutions as well. So it's being able to know how to procure them and recognise them as well.

LAWRIE ZION: Yeah. It seems all of those themes correlate with what we've seen in terms of what's happening in the media with deep fakes, et cetera, where there's already issues of trust with what's real and what's not in the media.

And there's a very good-- in the chat window. It's probably too long for me to just read out, but from Judith Bishop. I will try to just summarise a couple of the really important questions she's raising, because she draws attention to-- well, she asks what are our collective thoughts on the ethics of massive internet training data sets that have been used in training.

So again, how do we know if all the information swept up into that-- is this when we say your browser can accept all cookies? Are we giving data over without realising it, or is this data just being taken?

And Toby Walsh from the University of New South Wales has used the phrase "history's greatest heist" to characterise the hoovering up of knowledge and language from across the internet by global technology developers, such as OpenAI and Google, so getting back a bit to this--

And thanks for all the other points you've raised there, Judith, but maybe just to focus on that one for now about, is there a way that the responsible AI system, if it develops, as it evolves, can in any way step back or hand back data? In a sense, the same question we had with publishing when digital media got going. Is that going to be even technically possible?

RITA ARRIGO: Well, Data61, CSIRO are working on machine and learning that's currently in their field of purview. And I think it's really real. Like, your habits might change. And why do you have to be remembered for all those lollies and chips used to buy at the supermarket? You want to know that the world-- that your data is not being stored forever and those kind of things.

But the other area that I think is really interesting and something that we maybe learnt from social media is that when we started using social media, it was just so exciting to be able to connect with your friends and families and share stuff and all that kind of stuff. And then you realise that you were giving away a lot of. Your personal data. And I think that GPT is a little bit like that. It's an experiment. It's a tool we're learning from.

But if you're an organisation, I mean, maybe you're better off building your own foundation model and taking that foundation model and putting your own data in there and paying to use a service versus using a free service. And I think that's a lesson that we're definitely learning from social media, is that if you get something for free, there is often something that you have to give away. And that is something, a lesson that we need to keep in our minds, that if we were to--

And you'll notice that the South Australian government, they're always so innovative. They've just launched Edu Chat. It was on the news a couple of days ago. And they've actually taken this service from Azure, OpenAI. Obviously, Microsoft put a lot of money into OpenAI.

But now, they can take the foundation model and give it to you so that you can put your own data in. And so they've taken that, and they've made it a lot more secure, tested it. And they're trialling it. And I wonder if that is a real lesson that we need to learn around particularly generative AI, but also around these areas to take and suggest that we don't necessarily want to use this open content.

LAWRIE ZION: Yeah. Lots of questions coming in. And I am going to jump around a bit on a couple of these. Samantha Clifford's asked, "What approaches, Rita, are you taking at CSIRO," you, collectively, at CSIRO, "to bring the--" I assume, "to bring the AI principles into practise." And she's saying she'd love to hear more about this part of the work that's going on at CSIRO.

RITA ARRIGO: Yeah. So the first thing we worked on was that paper. It was super exciting to launch, and it's a great read. We obviously want to turn that into something that's very digestible and that is easy to understand. So we're going to do another webinar around it. And we're going to start to think about other ways that we can get this information out.

I think it's one of the-- the network is one where we want to learn from each other as well. So you'll probably see more case studies around how people are implementing that side of things. Also, more work around really demystifying what those principles mean and how do you put them into practise.

But that paper is a great one, if you're interested in downloading and having a read and then come along to some of the demystifying we're going to do in the webinars as well.

LAWRIE ZION: Just on that, Rita, because there's quite-- a follow-up question has emerged from what you were just saying, which is whether CSIRO's got any plans to implement some sort of certification process or program for the industry and startups. In other words, are you going to be able to, in a sense, endorse approaches that companies are taking through some consistent set of standards?

RITA ARRIGO: That could be a possibility. Yeah. And I think that would be an amazing value that we could add, because I think it would make it easier. And plus, we can really take responsible AI to the rest of the world. Because our solutions might have that capability in our region, and it could actually really help to grow an AI industry.

So yeah, we are grappling around that at the moment. The Ecosystem Directory was a super exciting initiative where we really wanted to create a way that we could find out about Australian AI companies. But going that next step would be super exciting.

And hence, we have a lot of help from the AI-- the think tank, which has currently moved into being the Responsible AI Think Tank. We have a heap of experts on that. Toby is on it and heaps of other experts. So we do look to their expertise. And the responsible AI network also has an advisory council. But we are open to heaps of ideas and heaps of feedback around that. But I actually think that would be a great idea.

LAWRIE ZION: Yeah. Another question from Sam Pincus which I'm glad has come up, because this event is being co-presented by two universities. "At what stage," he asked, "do you want"-- or she. "At what stage do you want to start intentionally training our students to become fluent or AI natives?" And I'd add to that, what would the training consist of if you did?

RITA ARRIGO: Wow. Well, you definitely have to-- I mean, ED chat in South Australia is for high school students. I was at an 18th the other day, and all they were talking about was how they used ChatGPT to plan their essay, or plan this, or plan that, or look this up, or do all that kind of stuff. So I don't know. Kids are going to be the first, I think, to attack it.

But I had a go at some of this stuff around the micro-skilling. I know that universities have got a lot of AI input into their syllabus at the moment, looking at that. But I think it would be really amazing for business schools to really look at a responsible AI or ethical AI kind of approach as well.

LAWRIE ZION: Yeah. It's a huge endeavour, collectively, across all these different areas. Ultimately, what's going to make this feasible is, it's in the hands of the government to support and to fund a lot of the basic, if you like, knowledge infrastructure that we're going to need in Australia to be able to manage and actually develop for the social good some of these principles in a way that they become embedded in the way that tech's developed using these, using AI or other innovative tech.

RITA ARRIGO: No. It's definitely a super important initiative.

So I think your question related to more-- how a university is going to do it and how are we going to teach our young people about AI? Yeah?

LAWRIE ZION: Someone's got to make the investment from the word go and say, we have to do this. We can't afford not to be doing this.

RITA ARRIGO: Mm. Mm. Yeah, I agree. It's really interesting because having had an AI education from a big tech company, it's very-- it's interesting to figure out, how do you take that thinking to market? Because often, students will not just want to learn theory. They'll want to put the theory into practise. So I know that micro-skilling can be done through working with cloud providers, et cetera, that have a lot of these tools available.

But see, I think there's also lots of different elements in educating ourselves in AI. It's not just the technical skills. It's also, how do you understand the user experience and the way that system works? And how do you implement a lot of those human societal impacts into that? And so it's that combination of technical skills as well.

And I think one of some of the biggest innovations have come from this multidisciplinary engagement when you bring a psychologist and AI together, or you bring manufacturing and AI together, you bring engineering and AI together, or you bring leadership and AI together. So I think it's not-- I think it's about embedding AI in all the aspects of learning.

LAWRIE ZION: Yeah. That's a very interesting observation. And I think one of the premises of this group technology by societies, we need to bring the various academic disciplines together to

collaborate on projects that will then have more value, we hope, to communities and to industry as well.

There are a couple of questions coming up. And you hear this a lot as a concern relating to the sudden emergence of some aspects of AI, where someone can do a Nick Cave song and it's not Nick Cave. Are we actually going to find ourselves as humans-- I'm paraphrasing these questions here-- locked out of the creative arts?

Because so many of the things that we do as musicians, or writers, or creators, and planning and writing essays, for that matter, will be possible to-- machines will be able to do this. So why would humans go along with this, or why would humans try to compete, I should say?

RITA ARRIGO: Mm. It's a really hard question, Lawrie. But I can tell you, I am involved in the art scene and the creative scene. And I get a real debate from people in that.

The evolution of creativity, when you have the ability to take away some of those-- writer's blocks or the grunt work that might be associated with having to have an incredible camera and not being able to do some of those things and opening it up for people so that you don't need months and months of understanding on how to use Adobe to build what you need.

So a lot of people are suggesting that this is actually going to be a bit of a creative renaissance, because you'll be able to really see the real creativity from the machine creativity and how different that is.

And I remember when ChatGPT first came out. The very first thing I did was put in the title of one of my husband's paintings, A Rabbit and Bananas. And it was really different, what the machine generated and what my husband had generated. So I just think that there's always going to be this higher type of creativity that's human and this kind of grunty creativity that's machine-oriented.

And I think a lot of people-- like, I was at SCREEN FOREVER recently. And I was with a lovely lady that does a lot of work. Ophelia does a lot of work around generating narratives. And it's really helped people to have much better screen results, because it can be tested and made sure that it's follows the right principles and those kind of things.

So I don't think it's-- I mean, personally, I don't think it's the death of creativity. I actually think it's a creative renaissance, where we really start to learn to see better and understand works and have that awakening around what's the difference between something that a machine created and something that a human created.

LAWRIE ZION: What about human diversity, though? I'm returning to at some points my earlier by Judith Bishop. And Judith points out, because data is the critical foundation of AI, how will the implementation of Australia's AI ethics principles ensure that Australia's AI data foundation, critically including the large language models we use, are going to be both ethically sourced and truly culturally and linguistically diverse.

LAWRIE ZION: Can you program --that kind of diversity into a system, or is it a case of-- I'm using a very crude example here, but you get into a lift-- notice I didn't say elevator and--

RITA ARRIGO: Yeah. [LAUGHS]

LAWRIE ZION: --voices of and American accent. So that's a pretty sort of trivial version of what I think is potentially a larger problem in that, do machines, in a sense, cancel out some of the most--

RITA ARRIGO: Well, I'll tell you--

LAWRIE ZION: Engaging parts of our system?

RITA ARRIGO: Yeah. Well, if we didn't have AI today, we'd be stuck in a lot more lifts. Because that was one of the first things that was being used to predict that a lift was going to fail. And so you'd fix it before you got stuck in it. So that's been happening. But I do--

It is really interesting, because you've just given this impetus to kind of-- a lot of people think voice synthesisation is wrong. Like, why do we need to voice emphasise, and why do we need to do that? But there are use cases when it's going to be really useful.

And your lift example is a perfect example. Like, I was at ACMI the other day. And they were saying, they normally have to pay really great voice artists a lot of money to come in and read stuff and create all the stuff. And now, if they have a small change, they can just use a little voice synthesiser to change that instead of having to drag the person back in and change that side of things. So there are all these productivity benefits that will come through that space.

But I do want to pick up on the actual question that Judith was asking, which was around AIs about data. And I understand that there are actually no large language models in Australia today. So when even when Edu Chat in South Australia has been deployed, that is all going into another data centre that's not in Australia.

And so if it's about data, it's also not just about the data that we put in there. It's also about how we lead it. Like, how do we decide where we're going to put it and who's going to look after it and who's going to invest in it. And so the leadership is also one of the key factors for us here in Australia, and it's not just about data. It is about what we choose to do with AI as well.

LAWRIE ZION: We're getting close to running out of time. But I'm going to be cheeky and throw one more question in before handing over to Sam to wrap up the event. It's from Anne Stevens.

And this is an issue that I think, really, really, a lot of smaller organisations in the not-for-profit sector are thinking about, how can they use AI to support their business systems or services so that we resolve some of the intractable problems in society that only the nonprofit sectors really work in? How do they get in at low cost and innovatively developing solutions, not just as consumers of AI delivered by big tech.

So that general question is such a thing possible for companies where-- can there be a low enough barrier to entry for small companies to benefit from these technologies?

RITA ARRIGO: If there's anything that goes from the trends that I've been seeing with people no longer needing to buy hardware and everything being in the cloud and being able to pay monthly for

things and those kind of things, then I do think the barrier to entry is going to be reduced because of that ability for you to have a software as a service and just tap into it for when you need it.

Like, I know when I was trying to do some copywriting, I used this-- I had to pay $40 a month to use Jasper AI, $40, American. But the results were amazing. And so you'll be able to pick and choose from this category, catalogues of different solutions that you might not have to spend $100,000 to implement, but you could actually get exactly what you need for that small fee.

So yeah, I do think that there is a lot of possibilities for not-for-profits to tap into these services that are going to be available.

LAWRIE ZION: This just is such a good session. I could easily go for another hour, but we'll have to leave it there, unfortunately, for questions.

Rita Arrigo, I want to thank you so much for coming in being part of this today. It's a fascinating space you're working in at the moment. And a lot of us are going to be connecting in to these questions a lot more. And that's one definite thing about what's happening over the near future. So thanks for being part of today. And I'll hand it back to Sam now to close the event. Sam, to you.

SAM WILSON: Yes. Thank you. I see lots of applause in the background, and I can only add mine to it as well.

Rita, thank you so much for taking the time today to give us this introduction just to this area. As Lawrie said, we could be here for the next three hours just scratching the surface. And we'd love to, perhaps, do it again at some stage as well.

And I would encourage everyone to not only follow NAIC on LinkedIn, but also follow Rita and stay abreast of what's happening there as well. The conversation is obviously live and evolving. And Rita, obviously, you're very vocal, and you shared a lot of thoughts here. So Rita's a great person to follow here.

I want just to thank Paul for setting up this webinar. And Paul, your stewardship is always so welcome here. You're always in very safe hands, and it's wonderful. So thank you, Paul, very much.

LAWRIE ZION: I've done a human thumbs-up there.

[LAUGHTER]

SAM WILSON: Nice seeing you in tech, Lawrie.

[LAUGHTER]

And Lawrie, thank you for leading the Q&A so well, as you always do. As you say, this is only the beginning of the conversation. And we really encourage you to contact Lawrie, or myself, or Diane to really follow your nose and keep the conversation going. There are many webinars and topics and research projects that lie ahead of us.

And Rita, your thoughts about how business schools can get involved is welcome news, indeed. So we'll follow those ideas up as well. And perhaps we can get some deep thoughts about what you see in that regard.

RITA ARRIGO: We need everyone to get involved. So we're super excited to continue the conversation. So thank you so much. And it's been such an honour to be part of this session today. So thank you to Lawrie and Sam for hosting me.

SAM WILSON: You're very welcome.

And now, audience, thank you so much for your great questions. Wonderful questions. And we look forward to seeing you all at our next webinar. So thank you, and have a great afternoon.

LAWRIE ZION: Bye, everyone.

PAUL LAVEY: Thanks, all. Thank you.

[AUDIO LOGO]

[END OF TRANSCRIPT]