Journal & Proceedings of the Royal Society of New South Wales, vol. 159, part 1, 2026, pp. 111–126. ISSN 0035-9173/26/010111-16

2025 Royal Society of New South Wales and Learned Academies Forum: "AI: The Hope and the Hype"

Panel Session 4: AI in Practice[1]

Mary-Anne Williams, UNSW Sydney

Ros Harvey, Australian Conservation Foundation

Dan Jermyn, Commonwealth Bank of Australia

Salah Sukkarieh, University of Sydney

Mary-Anne Williams

This panel is the Academy of Technological Sciences & Engineering (ATSE) panel. We are celebrating 50 years. Our panel will focus on AI in practice, and we have three extraordinary panellists to share their actionable insights into using AI in practice from robotics all the way to financial engineering, in a good way.[2]

Ros Harvey[3] was the founder of The Yield[4] and really focussed on how to help farmers make faster, better decisions. Dan Jermyn[5] is actually a closet physicist, just to give him a bit of credibility in this crowd. He's the chief decision officer at CBA — you might know that CBA was in the top five last year, and now the top three AI leaders globally. That's pretty extraordinary. Salah Sukkarieh[6] is our token roboticist. Does amazing stuff, also related to what Ros does in agronomy. The way this panel's going to run is that each panellist has five minutes and no slides, to introduce themselves and their perspective and point of view on AI. Then we'll settle into a few panel questions and will have 10 minutes of Q&A. Please start thinking about your questions early on because we want to fill that 10 minutes. So, Ros, you'd like to go first.

Ros Harvey

A quick bit of background about me and The Yield and AI. The Yield was a technology company I founded in 2014, focussing in agriculture. Basically, it was a very early AI business. We were focussing on microclimate weather predictions to use in agriculture to combine with agronomy knowledge. Later, we also grew it into doing microclimate predictions that farmers could use to make better decisions about how to grow their crops from irrigation, when to fertilise, when to harvest. Then we grew the product more into predicting yield as well. We did two things. I sold the business last year to Yamaha Agriculture, and at the same time they bought a robotics company. Salah and I have worked together for many years as well, so that we could stitch together the data and close the loop. So we're predicting

1 This is a transcript of a presentation, available at https://www.youtube.com/watch?v=jP_Rp6Aypno&t=20s

2 See https://research.unsw.edu.au/people/professor-mary-anne-williams for Mary-Anne. [Ed.]

3 See https://www.acf.org.au/about/our-organisation/board-council/acf-board [Ed.]

4 See https://www.theyield.com [Ed.]

5 See https://h2o.ai/ai-100/winners/dan-jermyn/ [Ed.]

6 See https://profiles.sydney.edu.au/salah.sukkarieh [Ed.]

the sort of things that you should do to grow crops more efficiently, and then predict the yield.

If you can add the data that robots actually see and what they do and you can automate, then you can close the loop with data, which allows you to create more powerful AI models. We were a very early innovator out of Tasmania. The Yield now holds the patents in Australia, US, Europe, Singapore — major markets for using AI to predict micro climates. It's a very foundational underlying technology. The patent application date was 2016, a long time ago.

I'm now president of the Australian Conservation Foundation, which is what I'm doing post-The Yield. I advise Yamaha strategically. I also have the right to take the underlying technology into other verticals, which I'm experimenting with.

I was asked what the big challenges of this were in the real world, of which there are very many, particularly in agriculture, which is a challenging sector, as Salah and I both know. I think the really headline thing I'd say is, look, CSIRO says right now, today, 80% of AI projects fail.[7] Some other folks say 95% fail. Why? People don't know what they're trying to achieve. What is the actual impact commercially? They don't have the data. I work in quantitative AI, not in general language models, and getting the data in the right format structure — so that it's actually AI-ready — is an enormous challenge. I don't care which sector it is from: finance, which is very digitally matured, through to aviation, through to agriculture, which is probably the least digitally enabled. It's very, very difficult.

The other challenge is the human side. A lot of organisations are just not ready. And when I say not ready, not only is it data, it's their people. At the end of the day, in my experience, the technology's the easy bit. The hard bit is the organisational and people bit. There's a lot of enthusiasm out there for AI, and there will be a lot of money spent, and there will be a lot of failed projects if we are not very, very careful about how we focus its on solving real problems with real people, solutions that are then capable of being adopted.

The final thing I want to say is we also need to be very mindful about the cost and the environmental impact. This comes from bitter experience in terms of the cost, because, at the moment, the industry is being really driven by the major vendors who have no interest whatsoever in reducing cost. Effectively what this means is that they're encouraging architectures and data-management strategies that are not optimising for cost-effectiveness.

I like to say that you can choose three things: you can choose flexibility, you can choose performance, or you can reduce cost when you're actually engineering. But what happens is the vendors don't care about cost. Flexibility is great because it drives up the compute and the amount of data that you're using, but it's very bad for customers and very bad for startup businesses.

It's also very bad for the environment because we know that every one of those ChatGPT queries draws two watts of power and half a litre of water in data-centre costs. This is going to have potentially a very significant impact on our energy use and our

7 Hajkowicz S (2025) Evaluating and prioritising artificial intelligence projects: A guide for better decision making and investment outcomes. CSIRO, June. https://www.csiro.au/en/research/technology-space/ai/How-to-choose-and-invest-in-the-right-AI-projects [Ed.]

net-zero transition. AEMO (the Australian Energy Market Operator) is predicting that up to 80% of our data use could be taken up with energy from data centres, which is a really big issue.[8]

These things are solvable because technology is human-mediated and we get to choose. We can choose: the way we architect, the way we design, the way we use technology that both solves for our community. I'm a techno optimist, I love data, I love technology, I love AI — but I think we've got to be very mindful in the way that we do it in the environmental, social, and economic benefits and impacts.

Mary-Anne Williams: Dan, there's a lot to dig in from what Ros was saying. I'm sure you're all realising that she is a sort of volcano ready to erupt.

Dan Jermyn

I'm from the Commonwealth Bank. We care a lot about responsible AI and trust and accuracy. I have to correct a few things. We're fourth, not third, and it's just in financial services.[9] I'm also not a closet physicist — I'm a failed physicist, but it's a shame, I was quite good for a while there. I'll talk about finance broadly. I am a representative of the Commonwealth Bank as well, and I was reminded of that, as I walked in casually, I didn't know where this place was. I didn't realise it was quite so grand. I feel terribly underdressed. Also, dangerously dressed, because if you can squint closely, this does say powered by CommBank AI. As soon as the first person who spoke to me realised I work for CommBank, they wanted to know a lot about a lot of things that I do not know an awful lot about.

I think it's important to understand one fundamental principle in AI: it is evolving so quickly. It's very, very complicated and it's very multifaceted. I've been working in it now for a very long time. I would say there is still a huge amount of unknowns about almost every aspect of it.

Notwithstanding that, maybe it's worth explaining a little bit about who I am, where I've come from, why I'm standing here today, and what I want to talk about. As you can probably tell from the exotic accent, I'm from old South Wales. I grew up in Cardiff, and I'm from a family where nobody had ever been to university. I did physics at university. I thought I'll do the hardest subject at the hardest college I could get into. I was doing quantum physics, and then I did a PhD, which I subsequently abandoned in computing to do a startup, a technology startup. That went quite well. We sold that.

I ended up in a bank — the Royal Bank of Scotland — and subsequently here in CommBank, where I've been for eight and a half years. Very pleased to say I'm a citizen of Australia now. Growing up in Danwood, never in a million years did I think, "Well, I'll end up in a bank." Actually, whilst I was being a failed physicist at university, I did an internship at an investment bank. Not for me ever again, that's not a place I want to spend any time. And yet, as I look back now where I am now, I feel an unbelievable

8 AEMO's updated forecasting methodology targets rapidly growing electricity loads, following industry consultation. 6 Aug 2025. https://www.aemo.com.au/newsroom/news-updates/aemos-updated-forecasting-methodology-targets-rapidly-growing-electricity-loads [Ed.]

9 See https://www.commbank.com.au/articles/newsroom/2025/10/commbank-among-best-banks-ai-maturity1.html [Ed.]

sense of pride and privilege that I got to work at CBA.

I think it's important to recognise that a bank's not necessarily the most obvious place to do AI, certainly not to do it quickly. It's not a glamorous place, it's not always a popular place (as I discovered with that first conversation), but it is an important place. Almost everybody in this room has a bank account — many of you have had mortgages, loans; you'll have insurance policies and so forth. It's a vital part of the fabric of society and it's vital that we encourage people who care about that society and that community to be a part of the conversation on AI.

When I joined eight and a half years ago, our CEO, Matt Comyn, who was then the retail CEO, had this idea of the customer engagement engine. The idea was that we would serve our customers well. We would just try and be better for them. It's apocryphally one of the most brief business cases in history. The stories have grown in terms of myth, but it's written on the back of a napkin: just do good stuff for customers and trust me.

But it's kind of true. I was never asked to do anything other than do good things for customers. And we thought about our multi-channel approach. We have customers in branch and a call centre, huge numbers of customers coming to us in digital. How do you stay relevant? How can you be helpful? It turned out that AI was a really critical part of that story because it allowed us to try to apply meaningful interactions at a scale that simply wouldn't be possible without that kind of technology. Over the course of that, we've strayed into all sorts of really complicated aspects of AI and how you get it to scale. We've made mistakes, we've made breakthroughs, but we've done some incredible things that simply would not have been possible without applying AI with a purpose-driven workforce. We've returned $1.3 billion to customers through our benefits. We've dramatically increased the speed at which we can respond to emerging bushfire activities. AI plays a critical role in all of these things.

Moreover, the way in which we're able to defend our customers from the rising scourge of fraud and scams is absolutely critically underpinned by AI. The bad actors are using it. We have to be in a place where our communities can be defended. I'm really proud that for all the difficulties that we have to work in this industry we succeed, and I want to maybe explore some more of that within the Q&A.

Mary-Anne Williams: Brilliant, Dan. It's great to have a leader like CBA in that industry, leading. And we've got little pockets everywhere, but we need to scale and franchise very fast. So it's good to have CBA leading the way.

Salah Sukkarieh

My name is Salah Sukkarieh from the University of Sydney. My background's in an area called field robotics, which is the application of robotics in outdoor environments. Just by the nature of being outdoors, the environment, biology, geology, and chemistry all change. The processes that we have to go through in developing these systems to work outdoors 24/7 in all weather conditions in different terrains is quite complex. Coming from the university, I generally focus on where we can push the boundaries on many of the algorithms and the sensors and the systems that we have to build. In the context of what we're talking about here, if we have to look at it from that lens, we think about it from a physical AI perspective. So how does AI work outdoors? How does AI work and adapt and change as the environment changes? There are lots of challenges around there.

My particular background expertise — or the area that I like to focus on — is environmental robotics. It's really in that extreme natural environment space that I work. We've done a lot of work in industry engagement with the use of robotics and AI and so forth with mining. Recently, we finished a project with the NSW Rural Fire Service on using AI and robotics in first response into bushfire scenarios,[10] and a lot of work in agriculture, both grains and wheat and horticulture and exploring those different domains; more recently, interplanetary systems as well. So rovers on planets and satellite systems, etc. So it's a wide range of activity.

When you look at robotics 20 years ago, the focus was: how do you take a platform that already exists and is operating, then add some computers on board, add some sensing on board, and make it work. You treated each one of those subsystems as individual units: the navigation system will tell you something; the control system will do something; and the sensing system will perceive something out in the environment.

Where robotics is going now, it's a lot more embedded. There's a lot more of that activity happening between the different subsystems. We're starting to realise that the way you move or control the robot affects what you sense and how you sense, and that becomes important. That quality of information then dictates how well you're going to control the robot as well, so there's that underlying principle. I think one of the key, really novel areas that's being pushed at the moment now is in agriculture. The focus there is on biology.

In one particular project, we're looking at phenotyping. In phenotyping, you're looking at millimetre accuracy changes day in, day out, on a platform that's bumping and moving around and slipping because of mud, so there's quite a lot of complexity around there.

In another project with the cattle industry —the grazing livestock industry — it's a much more complex scenario, where you're dealing with changes in pasture, but also moving animals. What we're trying to do with this system is couple the farmer's response and the agronomist's response into that. So you get a much more complex picture of, not just hard data and AI systems that can work with data, but also soft data — like large language models and how they interact with each other in that form.

The challenges that we're seeing as we go forward are lots and lots of data coming through, processing that information, trying to do it in real time — starting to ask questions as to whether or not you keep pushing down the AI path or whether there are other spreads. For example, optimal sampling in the environment may get you better returns than worrying about large collection of data and how that works within that operational context. These are some of the questions that we're asking.

Mary-Anne Williams: I'm sure that ending on that point — it's not lost on the audience — the connection between the importance of collecting the right data and

10 See https://www.nsw.gov.au/ministerial-releases/innovative-technology-supporting-faster-smarter-bushfire-response [Ed.]

the smart toilet that we were talking about in the last session, right? Now consider the toilet with arms and legs and robotic capability.

We've got to try and find some common ground here. We've got a few threads happening, but across all of these domains — robotics, digital agronomy and finance — where has AI really delivered value? If you could also identify one or two blockers that you had to get over. Where's the value and what are the blockers?

Dan Jermyn: There are lots of places where we've certainly added value. I've been at CommBank for eight and a half years, been doing AI for most of that time. Lots of different ways that we've spoken publicly about what we're doing with AI for our customers. If you fast-forward to today, one of the things that we've talked about recently that's really powerful is in terms of fraud and scams. I alluded to it earlier. Customer losses to fraud and scams for CBA — I think I'm right in saying —from a high in 2022 have been reduced by 76%. It's a whopping amount of money. There's still a lot of fraud and scams out there. We can't be complacent. We're dramatically reducing impacts on customers.

One of the things that we're most passionate about —actually one of the most interesting use cases I've seen recently — uses generative AI. A lot of the AI that we started with was predictive. Effectively the bank — all banks — are really largely predicated on prediction. You try and work out if you should loan to a customer and you try and work out the way you should set your price points and so forth. Even fraud detection is looking at anomalous behaviour and trying to work out if it's dodgy anomalous behaviour or just someone's doing something a little unusual.

In the generative AI use case —it's a great story for Australia — we partner with a local startup called Apate[11] to use synthetic agents for scam baiting effectively. You've all heard scam calls, people purporting to be from a large online retailer, let's say, and saying, "Somebody's been using your account" and trying to get you into giving them your details. What we've been using is 10,000 synthetic, fake CBA customers who now, using generative AI, can very believably in real time speak to a real scammer and keep them on the line for a long time. They think they're speaking to an actual CommBank customer that they're just about to defraud. Actually, they're wasting their time. The record is, I think, 52 minutes so far. It's very good. We're always trying to beat that. And in one simple sense, every second that guy spends on the phone to a non-real customer is the second that he's not spending on a real customer, really trying to dupe somebody.

Much more powerful is the intelligence that we gather about how these criminal gangs are organising themselves, how the scams evolve over time. This is something that's dramatically increased our ability to be reactive. That's an example of how we're able to use to today's technology, really advanced technology, predicated on a long history of building the muscle to build AI safely in a heavily regulated environment in a large corporate entity that wouldn't traditionally be thought of as a place that could do very rapidly adaptive technology.

11 See https://www.apate.ai [Ed.]

I think you mentioned hurdles. I would just say very briefly that none of this is supposed to be easy. We'd like it to be, but corporate transformations are always hard. I think there was nothing that we couldn't overcome, but there was nothing that wasn't challenging. The challenges persist today. You can't be complacent about it. The space still evolves. If I think about what did we do that was really important, it was around trust, governance, risk, transparency, and thinking long term about how AI would evolve and did we have the structures and the procedures and the tooling in place to give us confidence that we could do that at our scale.

Mary-Anne Williams: That's really great. As I listen to you, I'm thinking, okay, are you already developing the strategy for when the bad actors have their own synthetic AI attacks?

Dan Jermyn: That's already here. And the domains like cyber and fraud and scams are really at the front line of what's happening with AI.

Mary-Anne Williams: Ros, what do you think?

Ros Harvey: Trying to sell is actually a very good indicator of what actually has value. I originally came out of a university; there's a world of difference between a proof of concept and the brutal reality of actually running a commercial business, of having to hit revenue targets. With some confidence, I can tell you in agriculture, the single biggest value is reducing labour costs. That is the holy grail that everyone is looking for, and they're looking for Salah's robots. The problem is there are very few of them, particularly in horticulture. It's different in broad acres, but we both worked in horticulture. The technology is behind the demand.

The part of the business I came from — which is basically quantitative AI – there are two ways you look at it: something that reduces cost, and something that creates value. I'll use a very concrete example: yield prediction. Why does yield prediction matter? If you're running a big operation, say wine — and we had Treasury as one of our biggest customers, a multi-billion dollar listed company. We only worked with large companies because they're the only ones who had the data. The little farmers did not have the data for us to be able to access. So the value is massive. It's documented: thousands and thousands of dollars per hectare, because, if you don't get it to the winery at the right time, it goes off. You've got to pick it at the right time. It's this massive logistics problem of getting the right grey wine at the right time into the winery.

It's a massive labour cost. You've got to move people around to pick it and get it in. And then you've got to also deal with longer-term forecasting around what your blends are going to be to mix the market. Knowing up to 12 months in advance what the actual yield is going to be using AI at the block level, which we do very well, is really a lot of value. The only thing I could ever sell it on was, I was going to get rid of the cost of sampling, because what they do at the moment is they wander around and they pick, and they actually hold a frame out in the field and count flowers. They use a model developed by researchers, basically a mathematical model to project yield. Then they do it twice because it's expensive, so they can only do it twice.

We could sell doing it using AI: using satellite, weather and their historical data, we could do it for less than they could do it themselves, and then everything else was up side, and that is what sold it. The reason it's very important is when you try to actually get someone to take this technology on in the real world, if you say to a customer, "I'm going to create all this value for you," what happens is their employers push down on them and say, "Great, you're going to create all that value. I'm going to cut your budget because now you're going to actually have all this extra revenue and extra value and you're going to do it more cheaply." Not surprisingly, they're not very enthusiastic about budget-cutting. So it is really important when you go into the reality of doing this is to understand, particularly for a startup, how you do that.

I think it's different in large corporates: I'm now working in other logistics and travel industries where they can more flexibly deal with value creation. For example, reducing stress and things like that. But if you're in a startup business trying to sell an AI product, I think cost is the big driver.

Mary-Anne Williams: I think all universities really focus on innovation and entrepreneurship now. When I left university way back — I won't even tell you what decade — it was a hurdle to learn to program. It wasn't actually a course, but you couldn't get your degree without it. I think we're really moving very quickly to a world where you don't get any degree of any kind without some understanding of the process of innovation and all those things that you're saying. That's got to be good.

Ros Harvey: There's one thing I will say — I think in this audience in particular — it's a challenge managing researchers, because we worked a lot with UTS and I've worked with Pascal in the past with NICTA. Researchers want to push the boundaries of innovation, because that's what you're rewarded for. You get publications and all those things, but often the people who want impact, which is different, they just want one step better than where I am. And this creates a tension when you're trying to do both.

Salah Sukkarieh: I think universities want to create impact too — just a different form of it. I think if you look at robotics and how it has formed over the years, regardless of whether you're working in, say, a bushfire setting or an agriculture setting or environment rehabilitation setting, you throw these robots out there, and the first thing that was always worked on was just to automate the platform, just to make it work, go up and down, do whatever it might be. Where AI — maybe if I could just call it machine learning for a little bit at first — where machine learning really helped was allowing these platforms to now understand and perceive their environment, to understand it's a tree, this terrain I can drive over, something's moving in front of me that I need to avoid. That perception role — how machine learning allowed you to collect all that information, detect things, whether they're stationary or moving, and then determine how to navigate around — that became a big thing. Probably for maybe 10 years, there was a lot of effort and a lot of work in that space.

Over the last maybe five years, the value of machine learning is in the control space: how you control a platform optimally and accurately is really gaining a lot of momentum. What I mean by that is the task space of a robot can now open up a lot more. It can do a lot more things, because now you can learn the policies and the parameters that you want to go through.

Where we're coming to now is probably on two areas of high value. One is the interaction of those two things together — machine learning and perception, machine learning and planning control, and how they come together and work as a dynamic system. When you start to see those systems work and operate, you start to see what you think is intelligence in some way. It's doing something, it's reacting, it's behaving in a certain way, it's doing it optimally, etc., and it can adapt and change.

The other part is the interaction between the human and these robots — working with firefighters or working with the farmer or working with the conservationist and trying to say: if you've got this platform, it's like an extension of a Google search into the physical environment. These robots are going out, they're doing things. How do you interact in natural language? How does that process into real data? How does that come back to you as natural language? You get that synergy through that.

The challenge is to the second part of the question: when you throw into the natural environment, you miss communication systems. Connecting up to the cloud is not an easy thing to do.

How do you get these machine-learning models to actually operate on small embedded systems in the platforms, operate in real time, grab that data, and be able to do those tasks in real time as well?

Mary-Anne Williams: Riffing on your answer, I think it's great that you raise that interface between human and AI or robot. If we start thinking about agentic and autonomous systems or agents or robots, can we just outsource to them now just with some guardrails and a few trip wires? What are you comfortable actually letting these technologies do end-to-end or by themselves? Or what have you seen out there?

Dan Jermyn: There's a lot of ambiguity or what even is meant by agentic and/or autonomous. I can talk to some specific examples within the bank, which I think get the essence of what you're asking about: to what extent do you allow AI machine-learning systems to do things without intervention? I talked about the fact that we've been doing AI for a long time at CommBank. We're in a very heavily regulated industry. Everything we do, quite rightly, gets scrutinised very heavily. One of the fundamental principles around that is accountability. I'll explain a couple of systems in which things are happening autonomously in the way that you describe. It's important to be very, very clear that there is a human who's accountable for that system and checking that it works. If something goes wrong, it's that person who has to explain why, not the system.

One very obvious one is from the very first large-scale AI system that we introduced: the customer engagement engine. To very quickly recap: if you're a customer of the CommBank, and we're attempting to work out if you come into the branch, if you contact the call centre, or use the app, what's the best, most relevant thing that we could talk to you about? What are you likely to want from us? We're trying to be relevant. We can never get that 100% right, but we use AI to try to figure it out. There are hundreds of things the bank could be talking to you about or you may want to talk to us about. It could be: you want to change your credit card, you want to take out a home loan, you want to close your account. It could be anything.

We have several machine-learning algorithms trying to work out what you might want: the propensity and the likelihood. The customer might be eligible for lots of those conversations. The AI is trying to work out, of all the things that you're eligible for, which is the one that's most likely to be relevant. If we can get that right, more often than not, it's a better customer experience than just doing it randomly, which is what most companies would do. Those algorithms are running autonomously, they're self-learning, based on customer feedback, in real time. The very first was eight or nine years ago.

Mary-Anne Williams: Well, I don't know how you're going to get from four to one now if you're already doing that.

Dan Jermyn: I know. The bar is very high, but they're self-learning algorithms. We saw the power of that during COVID. We think a lot about the risks of AI because there are lots of risks associated with the technology that's evolving so quickly and are so complex. What we found in the implementation of this is no accident. It's by design: where we implement AI reduces our risks. At least that's what we're attempting to achieve. This is a great example. We have this autonomous self-learning system. We have lots and lots of automated checks and balances, by the way. From COVID, one of the principles underlying it is customer channel preference. We saw a couple of days when the model started to really drift in terms of expected behaviour. Some of that was to do with channel preference stuff. Alerts kicked in — bing, bing, bing, bing, something a bit wrong here. Maybe a data feed hadn't turned up. We defaulted to older models.

Mary-Anne Williams: Okay, Ros, what about you?

Ros Harvey: I think we probably are doing it customer-facing and then internal. The only way you can make this technology scale internally is to automate the process, AutoML, to actually create the models and results, because it's way too expensive with high labour costs, particularly the cost of data scientists. We have an automated ML (machine-learning) environment that is supervised, because obviously you need to have serious alerting set up on this and you need to be able to constantly measure your own accuracy of the models. We have detailed feature sets that we've developed through lots of contextual learning. We have auto ML environments that just trawl through different models, different combinations, to actually come up with new model results. That's internally focussed, if you like. We're increasingly using AI for coding internally.

On the external, I think we're basically at an automated system with automated supervision and accuracy checking in yield prediction and weather prediction. It's supervised in the sense that you've got an alerting system that tells a human if things are going off outside.

Mary-Anne Williams: What's the worst thing that can go wrong?

Ros Harvey: The worst thing? Well, we're not in an industry where someone dies if we get it wrong. What could go wrong? Well, for example, say you decide you have too many apples growing on a tree and you're trying to get the right number of apples on it for the season, so that you're not using unnecessary inputs that are going to give you too many little apples. Then you apply a chemical, a hormone, a product that effectively strips the crop to get you to the target number of apples. Then it has to be applied in a particular window, because if you don't get the temperature and humidity right, it's a little bit like napalming your tree. That is catastrophic. You can lose a whole season if you get it wrong. It is actually quite important to get the weather prediction for this thinning right, and you'd want to have a lot of confidence. You need to be careful on what your use cases are and what your risks are and what your confidence in your model are to do that. So, yes, you can have pretty bad commercial impacts.

Mary-Anne Williams: The idea that you can ahead of time specify how many apples you're going to allow this tree to grow?

Ros Harvey: Very known science. Developed by biologists.

Mary-Anne Williams: Okay. I have more respect for apples.

Ros Harvey: It's more about getting the weather prediction right, because our weather is on these big global forecasting systems, and rich countries like Australia add weather stations and we grid down. Maybe you're going to get to a 4 × 4 kilometre grid, where we use AI to predict inside the growing tunnel, inside the canopy, very, very specifically what the plant is experiencing. You can only do that if you're forecasting weather every 15 minutes, generating new predictions that are going out for every variable, massively out, up to 10 days and running this all through models. You have to automate this. You just cannot do this any other way.

That's where there are significant risks. What that means commercially is you might choose not to release a product that says, "Apply this chemical in this window." You might say, "What I'm going to tell you is what our best model estimates of these conditions are." It's a decision to support, not an automated system. With time, it will become automated, but I think we need these feedback loops so that you can effectively ground truth by what you're seeing.

Salah Sukkarieh: It was a good point that you made at the end, Ros. We used to apply AI systems or machine-learning systems, just black boxes. But the research over the last few years has done a couple of things. One is developing new machine-learning algorithms that also predict the level of certainty of their results. That has become an important element because now the output isn't just, "Here is the answer that I'm going to give you," but, "This is how confident I am about that answer." To follow on the decision support tools, what that also means is that if you keep running those algorithms, then what you give the user is options and different levels of risk maybe, or different levels of uncertainty.

The other part from the algorithm side has been recently how the machine-learning algorithm makes decisions. You're getting that output as well now, not just the final result, but: what it's doing, what it's processing, why it reprocesses something. That thought process, the algorithm itself goes through, helps the end user. Both of those situations have allowed the end user, the operator, to have a lot more confidence in the system for it to behave on its own. In our applications, we've been able to work from TRL[12] (Technology Readiness Level) 1 all the way through to TRL 9 in these different applications. What you find as you get further and further up the TRL is some of

12 Héder M (2017) From NASA to EU: the evolution of the TRL scale in Public Sector Innovation. *The Innovation Journal: The Public Sector Innovation Journal* 22(2). Available at: https://innovation.cc/document/2017-22-2-3-from-nasa-to-eu-the-evolution-of-the-trl-scale-in-public-sector-innovation/ [Ed.]

those machine-learning algorithms are not trusted by the end user for obvious reasons. A 400-ton truck on a mine site, a 20-ton fire truck, a 1-ton tractor, etc. You don't want these things running around spraying anywhere or going off-track. What you end up having is machine-learning algorithms and a change in operational behaviour. The two things are starting to come together. I used to operate in this manner, but now because of these machine-learning systems that are going to allow me to do something more cost effectively or repeatedly or whatever it might be, I will change the operation to allow that system to behave in that manner. So we're seeing those changes as well.

Mary-Anne Williams: I'm going to ask the next question with a little bit of a twist: we want to explore human-in-the loop adoption. Since you're all leaders and doing fantastic innovations, we want you to surprise us and give us an insight, an actionable one that we couldn't get out of ChatGPT. Something that is surprising because you've found it unintuitive when you've been working and building these systems. Just what has made you stop and think?

Ros Harvey: I've got one that I find difficult to believe. It comes going back to the day in oysters. We were predicting salinity levels. Salinity is a proxy for the runoff of water; because oysters are fresh filter animals, the health authorities say you can set a salinity level as a proxy for runoff by catchment. When they go over that salinity level, the growers stop harvesting. It's bad financially to do that. We started off by just trying to monitor real-time salinity. We did this with CSIRO, during the day, and that was very effective. Then the grower said, "Well, couldn't you tell us when it's going to actually be shut?" i.e. predict, which actually was an important innovation that we would predict. This was 2014. We had to come up with a machine-learning model to do this, to predict.

The thing that really surprised us, about after a year of doing this, is that we were taking data feeds like rain, wind, etc. And we were taking the live feed of the salinity and effectively building a model that would predict salinity from these environmental variables. The feed went down unexpectedly, but we continued to predict very successfully without the actual environmental variables. How can this be? I just didn't understand this because I was thinking about it as like an offset thing. And the team said, "Look, at the end of the day, the best predictor of something is itself. If you get enough data and you're feeding the error rates back into the model, it can actually predict itself." And I'm like, wow, that's amazing. We turned It off, but it continued to predict pretty accurately.

Mary-Anne Williams: Okay. Dan, what do you think?

Dan Jermyn: It's continually surprising. And that's both exciting and quite horrifying depending on precisely which job you have. In the example I was giving earlier, of the customer engagement and those models that were degrading, it turned out to be the earliest signs of people modifying their behaviour as a result of COVID. So the models weren't wrong. They were right. Our expectations were wrong. There was something happening at a macro level that the AI had cottoned onto that we hadn't. That was remarkable. That was not an intended consequence. That's not what the machine was for, but we were very pleased to see that it had "worked" and is accurately telling us what was happening out there.

More recently, the way people are using the tooling and what it means, I suppose, is it's evolving so quickly. If you had asked me a couple of years ago, we would have said a lot of what's happening in code generation AIs is very exciting.

It's going to democratise things, it's going to lower the bar, we're going to have a lot more people being a lot more productive with coding. What we've actually found as it got more and more sophisticated is the really elite engineers are being wildly more productive as a result, but the incremental benefits for an average coder are not really there in the same way. Rather than thinking about a world in which we're going to have lots of loss, more of these engineer specialists, actually where we're getting the real value now is from folks out in lines of business, maybe in more traditional areas, like risk management or in finance who are creative problem solvers and deep thinkers. A lot of the very, very best use cases that we get are being driven from areas that we definitely wouldn't have seen even a year or so ago.

Mary-Anne Williams: Salah, you have the final word.

Salah Sukkarieh: Maybe two examples very quickly, both of them related to how the human, the end user, has reacted to the output of the system. I'll take Ros's previous example around the apple industry. We used to work on perception systems on robots that could identify and count every single flower down a row, and then build a flower density map across a whole paddock or an apple farm. Six months later, you come back again and you do the same thing, detecting the individual apples. The idea is that this longitudinal study and data collection allow you to predict yield, so you keep going through that process. Really great farmers would stand on top of the hill, look down at the paddock, look at the approximate flower density, and estimate how much yield they would get at the end of it all. And we would go through the process and calculate the amount of yield as well, and we'd come within plus or minus a percent. So his estimates were perfect; that showed that the machine learning also could do it in real time and over that period of time, but you got roughly what the farmer was predicting by looking at the paddock. The difference was, we noticed that there were density changes across the different rows. So he was very good, repeatedly predicting the estimate of the yield because he would look at the flowers and he would then place the beehives for pollination purposes and get the yield that he was predicting. But you could then optimise where you placed the beehives to get a greater yield. That was not known up until that point, but there was a lot of reaction in terms of not wanting to do something. And that was one learning.

The other one was work that we did with Qantas on flight path prediction, fuel prediction, and trying to get optimal trajectories for the aircraft to fly, given the winds. That was another scenario where we were predicting behaviours of where the aircraft should fly, which looked completely different to the last 20 years of dispatching of flight plans. But once the upper management forced the dispatchers to fly the routes that they were going through, we were getting routes going right over Antarctica, for example, following through paths that you would never follow across to the US. They were getting 1%, 2% fuel savings on these flights, which is quite significant — another one where, from our perspective, we're applying machine learning and optimisation algorithms etc. to a problem.

We feel like it should do better, but then you've got— I don't want to call them the gatekeepers —the individuals who are so used to certain behaviours and processes; how they react to that result is quite novel.

Mary-Anne Williams: Okay. Well, I've got to apologise for the audience now because I've clearly got three optimists here. Well, all we got was serendipity. We didn't hear about all the disasters, but we're now going to turn to questions.

Q1: I'm an AI ethicist working in the federal government at the moment. I think this is probably for you, Dan. In 2019, our government published eight AI ethics principles,[13] and we have a body of work from, I think, Fifth Quadrant called the Responsible AI Maturity Index that looks at the relationship between effectively 500+ organizations' stated commitments to the eight principles versus how well they think they're operationalising those principles.[14] I just wanted to get a sense at CBA, how effectively do you think CBA is operationalising those AI ethics principles? What have been the main learnings from the challenge of attempting to do that in the real world?

Dan Jermyn: It's a great question. Yes, we're familiar with the Fifth Quadrant stuff and I've been through that process. I think it's a great piece of work. Actually, it's an involved answer because I'm very proud of the work that the team does on responsible AI. At the same time, we're very much grounded in the fact that we're working very quickly into a space that's evolving very rapidly. We need to be on our toes, and I think it behoves us all to have a certain humility. One of the things we've done is partner with a lot of very smart people externally, such as Bill Simpson-Young from the Gradient Institute.[15] There are other many great academic and other industry partners that we've very thoughtfully worked with, specifically because they have an independent and external lens on what we're doing. I think that's been really key. I certainly would like to think that we invite and receive — and quite rightly so — a lot of scrutiny. It's okay. It's good. Should be that way.

I think we're doing well. I certainly think that we've put a lot in our research agenda. We actually have a lot of really good muscle in this into matters of transparency, unintended bias detection. This is really what our work was predicated on from the very start. There's a lot of well-honed expertise at managing risk well. There's a lot of execution muscle in risk and governance within the bank naturally. What I think we could get better at is probably a good question, where we need to evolve.

Certainly, we're pushing very hard at innovation, at the boundaries of some of the things that are happening in, say, agentic AI, for example. We're very excited. It's actually an example where we have other autonomous parts of processes running and we're seeing better customer outcomes, but we have an ambitious agenda. We know that it becomes an ever more complex system as

13 DISR (2019, 2025) Australia's AI Ethics Principles. https://www.industry.gov.au/publications/australias-ai-ethics-principles [Ed.]

14 Fifth Quadrant, The Australian Responsible AI Index 2025. https://www.fifthquadrant.com.au/responsible-ai-index [Ed.]

15 https://www.gradientinstitute.org [Ed.]

it scales. The way that we're thinking about evaluation frameworks are the safeguards, the guardrail models that we've put in place for our customer-facing generative AI. This is something where we've done an enormous amount of work, but as the models get smarter and smarter, the work that we have to do to keep them safe becomes harder and harder too. So I am pleased with where we are so far. In fact, in that evident AI index, we were ranked number one in the world for responsible AI leadership, which is very good for Australia. Australia in general has a great tradition and a name globally for responsible AI.

Q2: If we'd had this meeting 50 years ago, we would be talking exactly the same way about the start of the digital revolution. What was interesting back then — because I actually started my career a bit before that — was that all the software had embedded assumptions. Of course, understanding the impact of these embedded assumptions was critical to understanding how you should use the software and in what particular case. Does that apply today in what we're talking about in terms of the embedded assumptions of all the software and tools that we use? Are these easily found and understood so that they can be used properly?

Dan Jermyn: If you want to gather your thoughts for something more erudite, I think it's a brilliant question. I think it's a more fiendish challenge with AI than it was in the digital revolution. I think it's implicit in all technological shifts, but I can give you a tangible example of some of the work that we've done with CommBank, because we worry about that. We worry about large foundational models that allow us to do incredible things, but they're trained on data. We don't always know what, we don't always know how, and therefore we don't always know exactly what the outputs may be. So how can we answer or at least investigate some of these questions? We did a study. One of the very best parts of my job is where we get to work with our behavioural scientists and the human-centred AI work. That's really not about the technology itself, but how people react around it and what it really means.

We did a study using large language models. Most, if not all, of you will be familiar with the late Daniel Kahneman and *Thinking Fast and Slow*[16] and Prospect Theory[17] and lots of well-published studies about what people will do if you say there's a 50% chance of you getting \$5 or a 1% chance of getting \$1,000: people will make "irrational" choices about that. And we thought, well, what LLMs will do, so we created a survey. We surveyed a thousand real Australians. We created a thousand personas like the real Australians, and we asked them the survey questions are for this very repeatable, very frequently run study that always gives roughly the same results.

The results were remarkable. So we ran it and the first question gave almost identical responses. The Australians answered it the same way everyone always answers it; the LLMs answered it the same way everyone always answers it. Then we asked question two, and I think this was a mistake. The team told me they did this deliberately, but they got the numbers wrong by a factor

16 Kahneman D (2011) *Thinking Fast and Slow*. NY: Farrar Straus and Giroux. [Ed.]

17 Kahneman D and Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–292. [Ed.]

of a thousand. The real Australian people answered in a wildly different way to what the traditional survey would say, but the LLMs answered in exactly the same way as the non-mistakey version would. It's a really visceral example of how easily we can be duped into thinking that, well, this is a reflection of society and the LLM is humanity talking back — well kind of, but also kind of not.

Salah Sukkarieh: I was just reflecting over the processes again. There's a lot more software being written now, so there's a lot of activity happening. You think you're on your favourite model and then next week there's a new model and you think, okay, we've got to try that one because, for whatever reason, we've got to keep up with the curve and get that going. But the same processes still lie from requirements all the way through; anything you write, you bring that through. When it comes to machine learning, usually we're asking questions such as what is the data that's being fed and that goes to that learnt model. We're a little bit more controlled, in the sense that we have that control to be able to do something like that. I think it's the LLMs where there's a lot more uncertainty.

Even in the context of the research of applying that — not so much doing research on LLMs, but using them in the research context – there are a lot of questions still there around: what data was it trained on and how can I use it in that context? That's probably where the greatest uncertainty is at the moment.

Mary-Anne Williams: Well, let's wrap up. But I'd like you each to identify one skill NextGen needs to thrive in this world we're rapidly entering.

Salah Sukkarieh: Co-designing.

Ros Harvey: Communication. That means listening as well.

Dan Jermyn: Creativity.

Mary-Anne Williams: Co-designing. Communication. Creativity. Well, there you have it. Thank you so much for this stupendous panel. You have been amazing. Thank you.