Launch Webinar: Cognitive Warfare, Biological Warfare and Workforce Impacts

Online 27 May 2026 4.00 pm Sydney time

TRANSCRIPT

Video recording here

PROF GLENN WITHERS (HOST)

Let me welcome folk to this webinar today. It's great to have some serious interest in a very interesting topic in the eminently topical AI field . And we've got some very interesting presentations to come from a very expert panel. So this is some creative engagement with emerging issues around AI and particularly, of course, links to cyber issues for us at the Cyber Institute.

I'd like to begin by acknowledging traditional owners of the lands on which we meet today and pay my respects to elders past and present. And our topic today is AI consequences for Australia. We're going to look particularly at strategic infrastructure, biological threats and workforce impacts. I'm the chair, I'm Glenn Withers from the Social Cyber Institute and ANU and Blended Learning International and various other things. And our tech manager is Lisa Materano, but our key contributors are the panelists. And that's Dr. Gary Waters, Dr. Brendan Walker-Munro and Prof. Greg Austin. I'll introduce them a little more in a moment, but let's get the procedural protocols straight so people know what our arrangements are for this one hour we've set aside. We plan to just have 10 minutes address from each panel member. That would leave us 30 minutes of questions.

As the hour comes to the end, we'll have a conclusion of thanks. And also at the end, you'll see I'm introducing you to some follow-up webinars that we're organising with our UK and US colleagues on AI impacts on workforce.

The first of them follows in about six days' time, and then about a month after that and a month after that, with the Australian parties organising the first webinar, and then the UK does the second and the US does the third. You'll see the link to get the details on that at the very end of this hour. So, if we could move on now to the panelists who are here featured.

Gary's our first speaker, and he's a distinguished fellow with the Social Cyber Institute. He's spent a very substantial amount of time in the Royal Australian Air Force, retiring as an

Air Commodore, then as a public servant in Defence. And then in the private sector as head of strategy for Jacobs Australia. He's now an independent strategy consultant and has leadership positions for organisations such as the Integrated Institute for Economic Research and the Critical Infrastructure Information Sharing and Analysis Centre. He and I have worked together too for the National Institute for Strategic Resilience.

Brendan Walker-Munro is a Distinguished Fellow with the Social Cyber Institute and an Associate Professor in Law with the Business, Law and Arts Faculty at Southern Cross University. He's had several appointments in law enforcement roles across diverse government agencies. and is admitted to practice law in the Supreme Court of Queensland. He is a member of both the Queensland Council of Conduct Tribunal and the Disability Panel of CPA Australia. So bringing a law lens to our topic today.

Greg Austin is co-founder of the Social Cyber Institute, held various posts with the International Institute for Strategic Studies, International Crisis Group, East West Centre. He was a Professor of Cyber Security Strategy and Diplomacy with the University of New South Wales, Canberra and has had various public service roles including postings in Hong Kong, a parliamentary committee, and ministerial staff.

So, we've got a very excellent group to help us address this topic and drawing on our Social Cyber Institute Fellowship, for that is a great way to get going on this topic. So now over to Gary, please, for an opening presentation and looking especially at critical infrastructure. impacts of AI. So, Gary, over to you. Thank you.

DR GARY WATERS

It's always a pleasure to participate in these SCI activities. I get a lot out of them. Usually, Greg forces me to think deeply about what I'm going to talk about because he changes the goalposts around as we think about it. And I'm going to start actually with the SCI position or the proposition that Australia should treat AI or artificial intelligence as critical strategic infrastructure.

So I'd like to explain why I support that, especially as agentic AI could significantly improve cyber defence of our critical infrastructure. But I also want to point out what I think it will take to deliver AI safely and also to highlight the emerging social and national security risks if it is misused or poorly governed. So first to the use of agentic AI in critical infrastructure.

Agentic AI can identify cyber attack patterns, detect malware, and surface unusual network activity by supporting faster, better informed, offensive decisions. Seeing a network through the eyes of AI can reveal relationships between attack patterns, traffic and target systems. Agents can use this information to build a knowledge graph of the network to strengthen protection.

Critical functions are usually based on a complex web of dependencies and interdependencies, which are managed largely through digital means. It's impossible for humans alone to understand those complexities and potential attack points. AI will allow for increases in effectiveness and efficiency of risk monitoring and ultimately risk

mitigation and system resilience.

The second point I wanted to make is that there are implementation considerations. AI agents can be built quickly, but validating them takes time. User acceptance testing and end user engagement have always mattered in IT and automation. With agents, they're central to delivery. In each scenario, teams will need to confirm the agent understood the task, acted appropriately in context, and achieved the intended outcome.

AI is incredibly flexible and highly autonomous. but a default to building massive complex agents might be a step too far. For example, it might be better to let agents interpret context and make decisions and let technologies like RPA (robotic process automation), handle the calculations and rules-based logic. This separation of duties would then create more clarity, stability, and flexibility. And there is a cost. Even a simple agent can trigger hundreds of large language model calls in seconds. An automation program that understands consumption costs or consumption across both the platform and the LLMs can save thousands of dollars a month. Agents are only as good as their prompts their context and the version of the model they sit on when that model updates or new scenarios appear behavior can change immediately thus prompt engineering or engineering of the prompts and ongoing fine tuning are essential. Fine tuning based on real usage will improve accuracy, consistency, and alignment with intent. So essentially what I'm saying is this is not a tick and forget model.

The third point is that these delivery realities lead directly to governance, which is a topic I've done quite a bit of work on. Agents are powerful, but they can be unpredictable. They may improvise, hallucinate, or present false information very confidently. Governance is therefore essential, including controlled agency and human in the loop checkpoints to ensure work is accurate and compliant with policy and regulation. So for high-risk business- critical processes, action should be routed to an end user for final review and approval in all cases. For lower-risk work, agents can be granted more autonomy with clear rules and constraints. Without strong governance, agentic work will be difficult to trust, adopt or scale.

My fourth point is actually moving the discussion away from artificial intelligence to augmented intelligence. Artificial intelligence can erode the judgment and skills of those using it when people rely on AI output without evaluating it. However, if AI is designed to strengthen human understanding, supporting deliberation, context and choice, it can reinforce rather than weaken human judgment and skills. So to achieve that, we need to think more in terms of how agentic AI augments its human users. Augmented intelligence enhances human capability while preserving essential cognitive skills. It treats AI as a sophisticated tool that amplifies rather than replaces human thought. And in the think tank Glenn alluded to before, IIER. I work with my colleague, John Blackburn, on that. We will be working on a major paper on that for the next few months.

The fifth point is that there are social and national security risks that need to be addressed. Recent research shows that AI models can change people's political views in less than 10 minutes of conversation, adding to growing evidence that LLMs can be used as powerful tools of influence.

This raises serious concerns about their potential misuse for disinformation and manipulation of public opinion. This sits within cognitive warfare, the use of tactics that exploit human cognition to disrupt, undermine or modify decision-making by shaping perceptions and behaviour through available technological means. That's a definition we'll use in our IIER work. If poorly designed, AI use contributes to cognitive atrophy at scale. People may become less able to detect manipulation and resist persuasion. And that would make AI-enabled influence operations significantly more effective and therefore a major national security risk.

Now, I won't do it in the conversation, but if you want me to later, I can elaborate a bit more on cognitive warfare.

My conclusion about these points is that treating AI as strategic infrastructure requires more than simply capability development. It demands safe delivery practices, governance that builds trust, and a deliberate focus on augmentation over replacement.

Agentic AI can improve resilience in complex systems. but only if it is engineered, tested and monitored with the same seriousness applied to other high consequence infrastructure.

For Australia, the agenda is therefore dual. Build the capability and competitiveness while also reducing systemic risk. That means investing in delivery disciplines, validation, separation of duties and the management of cost and consumption, as I mentioned earlier. Establishing governance that matches the risk. and designing systems that keep humans meaningfully engaged in high-stakes decisions. At the same time, AI-enabled influence operations and cognitive warfare should be treated as a national security issue, not merely a content moderation problem. Planning, policy and institutional readiness need to assume that persuasion and manipulation can be automated and personalised at scale.

So that covers what I wanted to say about the proposition that Australia should treat AI as critical strategic infrastructure. But I really want to extend that a little bit further to suggest that the national security risk that I just discussed is a lot more complicated in that as frontier AI models, such as Anthropic's Mythos, become more powerful, it makes sense to me and to others that the countries in which they're built, such as the US and China, would want to preserve some control over sharing them. Therefore, Australia can't assume it will always have access to such models.

So we should start to plan accordingly, such as by negotiating government-to-government agreement, where the US might treat Australia as a privileged customer, as it does with other capabilities.

Until new AI models can be incorporated into cyber defence to better automate intrusion detection and response, and into programming to write cleaner software in the first place, these models, I suggest, will favour the attackers. A model that is powerful for good uses will tend to be also powerful for bad ones. And while safeguards can be built in, the

inescapable point is that as models get better, they'll become riskier even if they're not meant for malign purposes.

And we've seen evidence of that in the last few days. Australia should factor this into its planning by working through the US alliance to shore up. shore up some form of privileged access it is also worth pointing out that Home Affairs has now created the full-time counter-terrorism coordinator role to provide national leadership and coordination on counter-terrorism matters, including the prevention of and response to terrorism incidents. This role, now with Brendan Dowling appointed, and the attendant support he will receive will benefit greatly from use of frontier AI models. And I must also point out that AI will be vital in understanding the dimensions of economic security national civil preparedness, and integrated statecraft together with their dependencies and interdependencies with the military dimension. That was all espoused in Australia's 2026 National Defence Strategy, which broadened the concept of national defence.

And I would argue we can't even begin to understand what that concept means without AI. So in a national defence context, assured access to AI systems, the data they depend on, And the infrastructure they run on is not merely a technical matter, but a prerequisite for strategic self-reliance. While the National Defence Strategy of 26 acknowledged AI's transformative significance, it did not address sovereign AI in its full dimensions. What I would refer to as data access, data protection, and assured operational availability as discrete strategic requirements. That needs to be rectified. So in short, and as my closing comment, we need to advance AI responsibly. And as the CEO of NVIDIA has argued, that means doing so safely with thoughtful policies for wide adoption with everyone included.

It must be guided wisely so it lifts up as many people as possible. It will speed up human knowledge. It will solve once impossible problems. And it can help build a richer, stronger, more hopeful world. but we must be careful in doing so. Thank you.

PROF GLENN WITHERS

Thanks so much, Gary, for your five-plus insightful points. And we'll move on to Brendan, too, for his insightful points. Over to you, Brendan.

ASSOC. PROFESSOR BRENDAN WALKER-MUNRO

Thanks very much for the introduction, Glenn, and thanks for the presentation as well, Gary. Again, it's also great to be contributing to SCI.

I'm very pleased to be presenting from the lands of the Yugambeh people, Gold Coast and Brisbane, and pay my respects to any first Australians that we may have in the audience as well.

My presentation's kind of a bit more in the biosecurity lens, looking at the impact of generative AI. It's something that, as we're about to see, gets a lot of airtime. And just to sort of bring a little bit of reality or a little bit of a sense check to a lot of the claims that

have been made about, are we at one end of the spectrum? you know, world-ending bugs or we're at the other end where it's all just a little bit of hype.

So I'd like to open kind of just with a word warning. I've trained as a lawyer. I did do medical science at university way back when, but this is not a super technical presentation.

When I use words like AI, I know that this covers a huge number of technologies, a huge number of subfields. I use it as a catch-all. I'm not going to sit and sort of try to bargain with the definitions like most lawyers would want to. The same with life sciences. These are very, very broad terms for reasons that you're about to see. But the big one, the really key one is that these industries are just moving so fast. They're capturing up so much ground that I could come up with a definition today and it would probably be completely useless by the time we finish the webinar.

So what is it specifically that AI is bringing to the life sciences? There's a couple of things that this kind of raises. One is the disruption of traditional security mechanisms that we might have had in place around sensitive technology type research or around sensitive life sciences research. What you've got now is a situation where you have the ability to put this information into a system that might actually spit that information out later to someone else.

So you're putting information into a Gen AI or into the training data into an LLM, and it then gives your responses, your data to someone else. So that's a huge problem from not just a national security perspective, but a privacy perspective, a human ethics perspective, an intellectual property perspective. So we really need to be considering about how we're developing these models, the security that goes into them, what we're training them on, and what we're actually asking these things to do.

You'll see at the bottom of the screen is an excerpt from the Virology Capabilities Test paper that came out last year. Now, this speaks to point two, the lowering of barriers to entry for sort of non-technical actors of making potentially a biological weapon. You needed to do your degree, do a master's, do a PhD, go and work in a lab, have access to expensive machinery. Now with the internet and with LLMs, it's lowering those barriers to entry. So you need less and less of the basic grounding in order to be able to actually get access to a lot of this information. And that also speaks to point three, where we are sort of bypassing the traditional gatekeepers. So in an academic context, it might be the head of the lab who doesn't just teach you how to do basic biological or biosecurity stuff. They're also teaching you things like ethics. They're teaching you how to behave responsibly. They're teaching you all of those normative principles that sometimes we take for granted.

What the generative AI and LLM technology also brings is essentially anything that's iterative, anything that you would have had a lab tech do for hours, days, weeks, months on end, endlessly pipetting, you can now have a gen AI or an LLM system replicate that, do the modeling in minutes, seconds, hours. A paper from Google that came out last year with their alpha fold protein technology found approximately 2,000 new types of proteins in six hours.

If you had tried to do that body of work with human beings, it would have taken decades.

So we're seeing magnitudes faster processing for these types of developments. And last but not least, and that's sort of the thing that's really scaring everyone is We, as we sit here now in 2026, have never had a greater capability to create entirely synthetic life, things that have never existed on the planet before. And we now have far greater capability to actually create those things than we've ever had at any time in our history.

And, of course, that's kind of what's scaring people. We've got a lot of discussions around things like mirror life molecules that are generated with a different chirality [1]to the molecules that all carbon-based life on Earth is made of, and effectively saying, well, if you do it this way, our immune system won't recognise it, they won't react to it, and you can create a toxin that'll wipe us all out. We've had AI-generated bioweapons, things that you can have made to order, ask the AI to create for you a biological weapon and it'll spit you out the results. agentic ai Are your AI bots capable of giving advice on how to make a biological weapon? So not doing the grunt work for you, but they're more than happy to sit over your shoulder as the metaphorical digital devil and tell you exactly how to create these things.
And then, of course, it's not just biological weapons. It's more capable of trawling the internet, coming up with the same amount of data for making a nerve gas like sarin. So these are the sorts of headlines. This is what's capturing a lot of the discussion at the intersection between Gen AI and biosecurity. How realistic is any of this? Well, a lot of it actually is kind of hype. So there's a really big disconnection in the science between capacity. So the AI model could and capability.

The AI model does. So if you look back at the VCT, the virology capability test I mentioned earlier, it outperforms humans about 10% better than us.
So it could, again, capacity, it could potentially teach someone how to create a biological weapon, but they actually have to still have the knowledge, the equipment, the skills, and the grounding to be able to put that knowledge into use. The timescales on getting from capacity to capability are also wildly varying. Some of them say the AI could do this in a couple of months.

Some of them are saying, well, it could take decades and you might actually need a quantum computer to be able to do any of this. So the actual sort of timescales about when any of these threats might actually manifest are completely out the window. At the end of the day, dealing with a biological or chemical weapon You need to know what you're doing. If you don't, you're actually more likely to kill yourself and all the people around you than anyone you might be choosing as a target. So even just putting aside a biological weapon requires a delivery system, requires stability, it requires a lot of investment. Even just getting your hands on a pathogen that is dangerous requires a lot of commitment. It requires a lot of money, a lot of time and a heck of a lot of risk. And anybody who has got any sense won't be doing it unless there is a substantial amount of money that they think is in it. And even then, they're essentially putting their life on the line.

[1] A geometric property of asymmetry where an object or molecule cannot be superimposed on its mirror image.

Even if you were able to do all of these things, use the AI, make a biological weapon, get all of the bits and pieces, try to put them together in your house. A lot of the signs that you're making a biological weapon are really obvious, and they're going to put law enforcement, intelligence agencies, the AFP, ASIO on your tail very quickly. Case in point, what happened in Las Vegas? There was another lab that was also found in California. Now, what's really important to note is these labs were not people saying, oh, I think my neighbor's cooking up a bioweapon. What they actually thought they were doing was making meth. They called the police, the police came in, and it's only when they found vials labelled anthrax that they realised they had a problem. Full disclosure, I should say, again, this is a conversation article I've written. I'm not just shamelessly plugging my own work. But this was, again, a really clear signal that these people were doing something wrong. AI or not, actually doing this kind of work is going to get somebody's attention.

That said, I think there's three things, and I want to leave this sort of with the group to think about and to sort of talk about and discuss for the rest of the webinar. Three areas I think that we're going to start to see some real changes. The first is sequence security. So at the moment, you can go online, you can pick up a phone, you can order genetic sequences in the mail, have them shipped to you. Until you put them into a bacteria or virus, it's entirely legal for you to have them. Think about ordering the parts to a machine gun off eBay and leaving them unassembled on your kitchen counter. As long as you leave them unassembled, you're not breaking the law. So I think we're going to start to see some changes around that, whether that's through biosecurity conditions or export controls or even fit and proper screening. So obligations on sequence providers to say, hey, who are you actually sending this to? Those are things I think we're going to start to see.

The second, which is a real bugbear of mine, and I really hope comes to fruition, is the death of dual use. I hate the term. Dual use applies to everything. A toaster is dual use. If you don't believe me, Google toaster laser. Everything is a dual use technology. Everything has a potential military or nefarious outcome. I think instead we're going to stop looking at these technologies as dual use or good or bad, and we're going to stay okay. Who's making these things? Who's using them? What are they making them for? What are they using them for? The shift is going to be on those ends of the spectrum.

And then as part of that, I think we're also going to have this focus on behaviour. So just like building a lab in your house is going to make people think you're making meth, not bioweapons, people are going to be looking out for red flags of behaviour. And that might even extend to asking Gen AI providers, and they've done this in the US, to keep an eye on what is it that people are searching for? What are they using AI for? What are they typing in? What are they trying to jailbreak it to do? Are they asking it, hey, where do I find my local anthrax dealer? And if they are, what are you doing about that? So I think those are the three sort of areas that we're likely to see over the next coming years to sort of try and deal with this intersection of AI and biosecurity. That's it for me. Thanks very much, Glenn and Greg.

PROFESSOR GLENN WITHERS

Brilliant. Thanks very much for being disciplined and delivering on time, Brendan, and also very much for moving us into the bioscience area to complement quite often the dominant

engineering and IT type approaches. As a member of the Phenomics Australia Board, this is crucial. So I'm delighted to hear a bioscience approach. It's great. So, we'll move on now to Greg for a final panel discussion.

PROFESSOR GREG AUSTIN

Well, thank you, Glenn, and thank you, Brendan and Gary, for your presentations. Based on what we've heard today, I think we can invoke the decades-old statement, "Houston, we have a problem". Gary and Brendan have sketched out a range of interesting and challenging circumstances, demands and opportunities. And we clearly need to address those with an AI-ready workforce. I'm afraid the bad news is that Australia has moved on a fairly timid trajectory when it comes to developing its AI workforce, not least in the range of circumstances and conditions and roles which Gary and Brendan have alluded to.

Analysing workforces for AI is a highly complex problem, multi-layered, all sorts of perspectives we can take. We need to look at things like the balance between threat and opportunity. We need to look at levels of analysis. We need to look at adjacent policies like immigration, economic settings, urban planning. We need to look at the way in which the international market operates in terms of AI products and AI services. But most importantly, we have to first understand what's going on in our own country. We have a huge data deficit when it comes to understanding the AI-related workforce that it has.

So one of the determining elements of the workforce situation of Australia in artificial intelligence is shared by the rest of the world. We can talk about China's capabilities in AI as much as we like, but it's my very firm assessment and the assessment of others that the US sector is currently the pace setter on the global scene. Its government's capability is the single most advanced in the world, and the United States is the single country in the world with very high impact on global education and skills for AI. That can be documented in all sorts of ways, but that's certainly something that we have to take as a foundation argument.

On the one hand, it complicates our ability to set domestic priorities for workforce development and for education. But on the other hand, it gives us a great advantage. So if we're an ally of the United States, as thankfully for the moment we remain, we can piggyback on the United States, the United Kingdom, other allies to help address workforce deficits and to exploit opportunities for workforce development.

Without going into too much detail, I'd really like to emphasize that we see two vectors that are really important in workforce development for AI. One is the technical, which is about how we gain benefits and face new risks from use of AI systems. Benefits include speed and efficiency in defense, but vulnerabilities include enhanced exposure to enemy use of adversarial AI. Beyond the technical domain, we've got to look at the socio-political impacts. This is one of the most important vectors of the development of AI in Australian society and the Australian economy. And we need no better indication of that than the events in Rome in the last 24 hours, or more correctly, in the Vatican City, when the Pope, along with a senior leader of Anthropic, made very compelling arguments about the need to subjugate technology development to important principles of society, human-centric

philosophy and human rights. One of the phrases I remember most from the encyclical relates to the supremacy of human rights.

So we've got two vectors crisscrossing, some would say inseparable, and we can't talk about workforces in Australia for AI and workforce impacts without understanding the balance between the two. In fact, there's an interesting study being done by a commercial organisation called Awarded Tenders, which documents quite convincingly the growth in the Australian AI workforce (in the Australian Public Service) in favour of addressing socio-political impacts compared with addressing technical impacts.

So in addition to the two vectors, we've got to consider four levels of analysis when looking at workforce development in any country on any subject, really, but it's especially the case in respect of artificial intelligence. It's no good talking about what we're doing at the national level if we don't understand the impacts of the international political economy, if we don't understand the individual choices that people can make about where they work and how they work. For example, citizens in Australia, people living in Australia can work for US multinational corporations in AI without leaving home. They can do work from home. They can do work from any country in the world, by and large.

And it's pretty unrevealed, in fact, just how the various nationalities actually sit in the workforces of the major AI corporations. I could only find one major corporation, which actually revealed the nationality makeup of its workforce. And interestingly, after the US as the dominant contributor to that company's workforce, Israel and India were the two next most important contributors, not the UK and not any of the traditional science and technology giants of Europe.

So I'll leave four levels. I think you'll all grasp that. To analyze workforce developments in our project with our British and American counterparts, we developed three scenarios. One is conservative, one is moderate, and one is radical. And we can express them in terms of the level of appetite for reform of the workforce. Conservative being a low appetite for reform, moderate being variable and radical being a high appetite for reform. And to illustrate the last point, in 2021, the US National Security Commission for Artificial Intelligence made a radical proposal for reform of the United States AI workforce by setting up. an AI corps, putting people on scholarships effectively in mid-career to go away and study AI for a year and come back to help move the United States workforce along and to achieve the best effects that could be achieved from utilisation of AI, but also to prepare the best defences against enemy use of AI and to help sustain US global position in the world. So that's radical appetite for reform where you do radical things and everything looks different.

The interesting assessment coming out of this work, which we'll return to in a minute, is that Australia sits, sadly, somewhere between conservative and moderate. So we're using those three scenarios to try and understand how best to assess the trajectory of workforce reform in different countries.

I've got two slides here from this study I mentioned before by the company called Awarded Tenders. It has data from 2020 to the first quarter of 2026. It's a study based on advertised

vacancies in the Australian public service through the Commonwealth Gazette. And it indirectly highlights, that in comparison with the public sector, the private sector in Australia is very non-transparent when it comes to revealing how many of its employees are AI capable.

And the private sector also appears to be moving a little bit more slowly across the board than the Australian government. The figures here from Awarded Tenders based on the advertised roles in government vacancies are quite illuminating. We see defence and security as the dominant sector. So number one agency, Department of Defence, number three agency, Home Affairs, number four agency, ASD, and then number 10 agency, ASIO. So it's really an international and domestic security focus that has grabbed the Australian government's attention more when it comes to AI reform.

The Australian government has undertaken some very interesting reforms in staffing the AI function in the Australian public service. For example, they mandated the establishment in each department by July this year of something called a chief artificial intelligence officer and the associated roles of AI accountable officer.

These are radical moves, but we've yet to see, in fact, the end results of that.

One thing to notice from this particular slide is that the numbers of roles advertised outside of the security agencies are absolutely minuscule, given the scale of the opportunity and the scale of the threats. We should bear in mind that government recruitment of AI specialists is taking place through a mixed mode. On the one hand, we have recruiting AI specialists themselves.

On the other hand, we have contracting out to firms who deliver the AI talent without any public revelations, for the most case, of what's involved in staff roles and the numbers involved. But by and large, the Australian government's workforce response over six years for the AI opportunities is very small when it comes to overall numbers.

So in our work on AI workforces, we've paid some attention to how important the particular subfields of AI are. And we take the view that if you want to analyse the workforce, you've got to know what the flow is in the various subfields of AI. And what we see in this information, again from Awarded Tenders and their data set on jobs in the Australian Public Service, we see some interesting trends. For example, computer vision dropped from 8.7% in 2020 as a share of the number of ads to 1.7% in 2026. But we're seeing stability in some other areas.

The most important point here is the one I mentioned before, but chance is that the share of roles taken by for regulation or regulatory functions in respect of AI has surged from 2.9% in 2020 to 34.6%. So these are really important. reference points.

So to conclude, Australia is on a conservative to moderate path. The Australian government's policy can only at best be described as timid. We can discuss why that is and what are the consequences of it, but we have timidity ruling in Australian AI workforce policy. As I mentioned, we're strong on high-level AI and digital workforce framing. But

we're relatively underdeveloped in explicitly articulating AI workforce development.

We've got all sorts of nice motherhood statements and ticker box type policies in place, but we really don't have the depth and the granularity that we need for an advanced AI workforce.

One of the reasons is that we rarely get beyond the concept of general AI capability down to specific AI subfields. We do need to do that. We need much more attention to this and we need to find a mechanism to move beyond finding ourselves halfway between conservative and moderate, at least achieving some sort of moderate trend and reaching for radical. Let me stop there.

Thank you, Glenn.

PROFESSOR GLENN WITHERS

Thanks very much, Greg. A very helpful insight into the Australian Commonwealth developments in particular and workforce hires.

Q&A not available in the published Transcript.