ASPI
AUSTRALIAN STRATEGIC POLICY INSTITUTE

Submission to the Joint Select Committee inquiry on Artificial Intelligence

David Wroe (Senior Fellow, ASPI)
Justin Bassi (Executive Director, ASPI)
Olivia Nelson (Chief of Staff, ASPI)

September 2026

Note: The views expressed in this submission are those of the authors and do not reflect a single ASPI perspective.

Introduction

Artificial intelligence is the most consequential technology humanity has ever invented and likely will ever invent. We are on the cusp of creating machine intelligence that is more capable across the board than human intelligence. No country, including Australia, is ready for this historic transition despite the consequences for our economy, society, security and personal lives. The recent burst of attention and discussion is welcome, but we still need greater policy focus and public debate. This committee is a much-needed addition to those efforts.

Australia has two broad objectives.

First, we must capture a substantial portion of the benefits by establishing and growing a durable part of the global value chain here. The federal government's proposal to do so by attracting AI computing infrastructure investment and using that foundation to build a higher value industry at the model and application layers is a sound strategy.

The challenge now is in the execution, which includes building public trust by clearly explaining benefits while proactively managing risks to give Australians confidence that they are not unreasonably exposed to higher power prices, rogue AI agent cyber attacks or the many legitimate concerns in between. The government and Parliament should not wait for public trust before acting. This will only result in missed opportunities and deepen the existing anxieties about AI. Rather, as a core national security priority, leadership must nurture the social licence that will be needed in the long term. The end goal is to strengthen Australian AI sovereignty—a much discussed term that has different interpretations but boils down to ensuring we keep our agency and strategic options by controlling enough of the core inputs that we cannot be denied or coerced. This doesn't mean doing everything here but it does mean doing enough that we can trade on others' dependencies on us.

Second, we must protect ourselves by strengthening our defences and resilience at home, and helping to shape global governance to reduce the risks of uncontrolled and misaligned AI. The risks posed by superintelligence or powerful general intelligence are very real and must not be dismissed. The stunning pace of AI development, with growth both in the size and power of models, and in the autonomy and independence of agents, is outstripping researchers' ability to ensure

the technology is safe. The commercial and geopolitical imperatives to keep racing are forceful but not insurmountable. Australia as a credible middle power can and should express its concerns and contribute to solutions.

The first and second objectives are complementary because a deeper and more indispensable role for Australia in AI development will also give us a stronger say in the standards, rules and norms that can make it safer. The best way to protect the nation is to be an active player in AI's safe development. While our island nation provides many natural defences, it can't keep out AI risks unless we unplug completely. To be clear, AI development should continue because, managed wisely, it could be one of the best things that ever happens to humanity, making us wealthier, healthier and more fulfilled.

This Select Committee is an invaluable input to the ongoing efforts of federal policymakers and a vehicle for Australians to contribute to the direction of AI development. Its timing and purpose could not be better and therefore we strongly support your work. Given ASPI's policy remit, we are focussing this submission on the strategic policy and national security implications of AI.

Recommendations

1. Use Australian diplomacy to push the US and China towards mandatory frontier-AI safeguards, including safeguards calibrated to increasing model capability ("pacing the frontier").
2. Increase funding for Australia's AI Safety Institute, enabling it to recruit leading experts and contribute more substantially to global safety and alignment research. Particular capabilities should include independent frontier model evaluation; cyber/bio capability assessment; incident analysis; international information sharing; and access to models before deployment.
3. Work with other democracies to lead AI capability and therefore democratic principles of technology governance diffuse globally.
4. Implement the government's data-centre standards rigorously, especially requirements for providers to support Australian research, supply affordable compute, build local skills and invest in Australian supply chains.
5. Bolster programs to attract, develop and retain AI talent, treating skills and expertise as a particularly durable investment given uncertainty about AI's development.
6. Increase federal public investment in AI, including consideration of dedicated sovereign compute and model-development capability.
7. Use Australian data-centre investment as leverage to secure access to advanced US frontier models, particularly for national security, critical

infrastructure and strategically important industries.

8. Fund Australian research into defences against AI-enabled biological threats, drawing on Australia's strengths in medical and biological sciences.
9. Keep mandatory cyber-security standards—especially for critical infrastructure—continually updated as AI capabilities evolve, including through the SOCI Act framework.

Growing Australia's AI industry

Australia's immediate advantage is in building AI data centres (or "compute")—a fact that is widely established and agreed. This investment is important for two principal reasons: first as a foundation for a higher value AI industry that will strengthen our sovereignty and economic security, and second as a basis of leverage by which Australia can negotiate for access to the most advanced technology such as frontier models, primarily from the United States.

On the first rationale, the government's five "expectations" and planned national standards are a strong basis for encouraging this investment in a way that matches our national interests. We won't traverse all of the issues here but rather will highlight the importance of expectation five, which calls on big data centre providers to contribute to local research and innovation by supplying cheap compute to Australian companies and researchers, deploying skilled workers here to build the local talent pool, and investing in Australian supply chains.

These are important provisions and must be implemented precisely to deliver the economic security and sovereignty Australia seeks. We recommend the committee explore in detail with the government and industry how the government will ensure the national standards translate reliably into the desired national interest outcomes.

Those outcomes are admittedly hard to define because the precise direction of AI development is hard to predict, even if we know more generally that the technology is extremely powerful and important. The effort needed to develop an Australian frontier model company that could compete with the world leaders, for instance, would be enormous, perhaps prohibitive. Anton Leicht, from the Carnegie Endowment for International Peace, has estimated the cost of building a frontier company from scratch at US$500 billion. But Australia can and should grow its model development capacity at a more modest and realistic level, including by taking off-the-shelf open weight models and fine-tuning and enhancing them for our own bespoke purposes. One note of caution on this ambition is that many of the best open weight models currently are made by Chinese companies, which pose security and strategic risks and would need to be considered very carefully by

Australian users. Certainly for national security and critical infrastructure uses, they should be considered too risky.

Given the policy uncertainty owing to the difficulty in predicting exactly where the technology will go, the safest investment will be in skills and expertise. Programs that attract and retain talent should be bolstered. However the technology evolves over the coming years, Australia won't regret building a strong and skilled AI workforce.

The federal government should also consider more public funding for AI development. International comparisons are a little fraught but a fair estimate is that over the next five years, Australia is committed to spending $500m to $700m in public money on AI development. This includes no money for either sovereign compute or model development. Arguably, the anticipated private capital for compute in Australia with deals that guarantee reserved compute for Australian companies and researchers might be a substitute, but until those deals are in place, it is hard to be sure. Certainly nothing can be as certain as publicly funded and dedicated sovereign compute.

Comparable countries typically spend public money on sovereign compute or other key parts of the AI stack. Canada, Britain, Singapore, France and Germany have all committed multiple billions of dollars over the next five years to AI programs, mostly to compute, research, model development and talent. South Korea and Japan have earmarked tens of billions of dollars to compute, models, chips and more over the same period.

On the second rationale, building data centres that can be used to train new frontier models should give Australia leverage to negotiate access to those models. This access is not otherwise assured, even from Australia's ally the United States. The constraints placed on the release of Anthropic's Mythos Preview in April were the first concrete demonstration that both frontier companies and the US government will increasingly find it hard to justify sharing access to their most powerful new models. From their point of view, they are building systems that in the wrong hands could constitute superweapons and it would be irrational to allow permissive access.

If infrastructure in Australia is necessary to these models' development, the Australian Government has a much stronger hand in negotiating access, with all the understood safeguards and with Australia's reputation as a responsible country with strong laws and elite national security agencies. Australia should pursue the best possible capabilities for paramount needs such as national security, critical

infrastructure and industries of strategic importance, even if sub-fronter models are sufficient for many other uses.

Australia's current copyright regime is an obstacle to AI model training for reasons well canvassed in public. We urge all parties to reach a reasonable compromise as soon as possible given that the window of opportunity for Australia to become the natural second home for US companies to train models is finite. We are neutral on the precise mechanism to resolve the copyright issue provided it does not deter investment. Our only overarching observation is that a compromise involving updates to the regime is warranted given the emergence of large scale AI training raises novel questions that the existing copyright framework was not designed to address.

Strengthening our defences and resilience

AI is dramatically lowering the skill barriers to sophisticated cyber attacks, with some experts warning that capabilities previously confined to nation-states are now within reach of individuals with only modest skills. Again these threats are well documented in the public domain.

The Australian Signals Directorate is doing a superb job of managing the cyber threat transformation and we encourage the Committee to highlight the Directorate's work to ensure the private sector and public are aware of the resources available. National preparedness and resilience are the job of every Australian.

As AI systems become more powerful, the range of threats they can pose, either through deliberate misuse by humans or through misaligned and uncontrolled systems themselves, will continue to grow. New bioweapons created through the convergence of AI and synthetic biology are considered the next emerging risk.

Some commentators argue that AI companies should release frontier capabilities as quickly as possible because the more capability in the most hands will enable defence to quickly outmatch offence. This is very risky in the short-to-medium term because the sheer scale of existing and undiscovered vulnerabilities in the world's software will advantage attackers for at least a year or two until vulnerabilities can be patched. There are many fish in the barrel for malicious cyber actors to shoot. Over time, once vulnerabilities are patched and AI-enabled cyber defences are widely deployed, the offence-defence balance should even out and might even begin to advantage defenders. However, this strategic logic will likely not hold for

bioweapons and other future threats. Biodefences could involve inoculating humans against novel pathogens—a logistical, political and social nightmare compared with patching software vulnerabilities. Australia should therefore be very wary of arguments in favour of rapid frontier AI release and should use diplomatic channels to urge caution wherever possible.

In terms of federal responses, the government should continue to monitor the appropriate level of resourcing for agencies such as the ASD to ensure they can keep pace with the progress of the technology and the scale of its impacts. Given Australia's prowess in medical and biological sciences research, we would also urge the government to consider where it can direct public funding towards organisations equipped to carry out research into protections against future AI-enabled biothreats.

In terms of regulation, Parliament should focus on where it can make a difference in the areas of higher risk. The more extreme scenarios of global catastrophe—which again should not be dismissed—are obviously hard to legislate against. One clear priority is to keep mandatory minimum defensive standards up to date, particularly for critical infrastructure. We therefore support fully the independent review into the Security of Critical Infrastructure Act 2018 and the ongoing implementation of its findings. And we encourage further regular updates as needed, given the speed of progress in AI.

Reducing the global threat level is considerably more difficult. To take OpenAI-Hugging Face attack revealed in July as an example, it is difficult for any Australian legislation to prevent such a rogue agent attack had it been directed at an Australian company or entity. One consideration for the government and Parliament is to review legal liability laws so that an overseas-based AI company whose agents carried out a cyber attack in Australia without the company's direction or knowledge could be held accountable if their behaviour were deemed reckless. Given such laws would need to apply extra-territorially, this is a tricky area that would need to be considered carefully, noting that even under current US law, OpenAI doesn't appear to have committed any crime and Hugging Face has no stated plans to pursue damages or legal accountability.

More generally, the Commonwealth should continue working closely with states and territories to keep laws up to date to protect Australians—especially children and other vulnerable groups—from more immediate harms created by AI. Efforts to criminalise behaviours such as the creation and distribution of deepfake pornography and CSAM material are clear examples.

And it should contribute to democratic leadership of AI worldwide. The best technology tends to also set global standards, which reflect the values of the dominant nations.

Influencing global safety

For the more global and ultimately more catastrophic risks, Australia should deploy its international diplomatic influence as vigorously as possible to persuade the US and China—the only two countries building proper frontier capabilities—to cooperate on mandatory safeguards calibrated to capability development, what has become known in the US as "pacing the frontier". Neither the US nor Chinese government has endorsed any kind of slowing or pacing of frontier development, despite the risks described by many working at the frontier in the US industry, including the leaders of Anthropic and OpenAI themselves.

To enhance its diplomatic clout, Australia could grow its contribution to global safety research and development, notably through the AI Safety Institute. The AISI's funding should be increased to give it the resources to attract more leading experts and conduct more research. The world needs more safety and alignment research, and the countries that contribute to AI safety are also best placed to contribute their values.

Building trust

Public trust in AI in many countries is in a precarious state. A recent poll carried out by news outlet *Politico* found that 63 percent of Americans believe AI poses at least a "moderate risk" of destroying humanity. A Lowy Poll released in June found 64 percent of Australians believe the risks of AI outweigh the benefits. If it falls further, there is a real risk of a strong political turn against the technology.

Many of the concerns are legitimate, including the risk of catastrophe through misaligned or uncontrolled AI, and the likely economic disruption including job losses, at least in the short-to-medium term.

Governments need to earn trust in the development and deployment of the technology by persuading their citizens that they are responsibly confronting and managing the risks, as well as channelling the benefits for the good of everyone. The dismissive attitude shown by both the US and Chinese governments about serious safety concerns being expressed by the US industry and many independent experts is a problem over which the Australian government has limited influence, but it should nonetheless advocate to those governments and more widely to other

countries for greater global cooperation on safety.

The major risk for Australia is that falling trust in this country leads to staunch opposition to data centre construction. Given AI infrastructure is Australia's best ticket to economic development and global influence over governance and safety, the government should remain focussed on winning and maintaining public confidence on infrastructure building. It would be detrimental to Australia's national interest if anxiety and anger towards AI generally were channeled into obstructive opposition to data centres that would actually help Australia influence the technology's direction. Worse still would be the politicisation and populist exploitation of resentment towards data centres.

The government's commendable emphasis on building social licence should continue and, if necessary, be strengthened. And the Coalition's equally commendable support for data centre construction should likewise continue.

Community backlashes against data centres are likely driven by many factors, including power prices, but might also reflect the fact that data centres as the most visible manifestations of AI are serving as lightning rods for broader, nebulous concerns such as fears of job losses, catastrophes, harm to children, social dislocation and cognitive deterioration. Therefore the Australian government should bolster its efforts to address all legitimate concerns, including by regulating where that can be effective.

People cannot be won over with words alone; trust is earned over time through demonstrated action.

Conclusion

Australia's strategic AI policy needs to work on many levels at once.

It must ensure Australia has access to the best capability.

It must contribute to global safety by impressing upon the US and Chinese governments that uncontrolled and misaligned AI is a global concern they must address.

It must find and plug gaps in domestic regulations to protect Australians at home.

And it must seize the huge opportunities to build and adopt AI.

If the control and alignment challenges can be overcome and the power of AI can be

harnessed, it will determine the global pecking order. As the US President Donald Trump said recently, "Whoever wins AI, wins." He wasn't the first to make the observation. As far back as 2017, Russian President Vladimir Putin said, "Whoever becomes the leader in this [AI] sphere will become the ruler of the world."

The world's turbulent geopolitics is unlikely to stabilise any time soon. It's going to be a difficult few decades ahead, so reducing our dependencies is strategically wise. Given the unprecedented power of AI, big bets are merited. Not all of them will pay off. But those that do will pay off handsomely and won't be regretted by future generations of Australians.