

Risks and Influence in the Age of AI

Submission to the Joint Select Committee on Artificial Intelligence

September 2026

Submitted

14 September 2026

About Good Ancestors

Good Ancestors is an Australian charity dedicated to improving the long-term future of humanity by providing rigorous, evidence-based, and practical policy recommendations for Australia's biggest challenges. We have been deeply engaged in the AI policy conversation since our creation, working with experts around the world and helping to organise Australians for AI Safety.

Contact

If you would like to discuss this submission, please let us know at contact@goodancestors.org.au.



**Good
Ancestors**

Table of Contents

Summary of key points.....	2
AI is more capable, and consequential, than people realise.....	3
AI is moving fast.....	3
Greater capability brings greater risk.....	4
The harms we see today are warning shots.....	5
Australia needs influence in the age of AI.....	6
Hosting frontier AI training gives Australia more of a say.....	7
Building inference capacity would boost our economy and security.....	9
Capturing the benefits of AI training and inference requires a copyright solution.....	10
Australia needs information and expertise to act on AI risks.....	13
Australia needs visibility into AI risks.....	13
Australia needs expertise to understand and respond to risks.....	13
Australia needs regulations fit for AI.....	14
Mandate gene synthesis screening to address AI-bio risk.....	15
Establish who is liable when AI agents cause harm.....	16
Conclusion.....	17
Attachment A: Building AI while Compensating Creators.....	18
Attachment B: What We've Heard So Far.....	28

Risks and Influence in the Age of AI

Risks from AI are increasing alongside capability. Australia has little influence over a technology on track to upend geopolitics, the global economy, and national security. To manage these risks and protect our national interests, Australia needs more influence over how AI develops, the right information and expertise, and laws fit for the technology.

Summary of key points

AI is more capable, and consequential, than many people realise. Advanced AI can be deliberately misused, it can malfunction (including escaping human control), and it can cause systemic harms such as labour-market disruption and the concentration of power. These harms are already occurring, and are warning shots of worse to come. Australia is exposed to these risks, and unprepared for them.

To manage these risks and protect national interests, Australia needs:

1. **Influence over how AI develops.** Australia must secure valuable positions in the AI value chain, and use them to secure our national interests.

Hosting frontier AI training may be the most valuable position Australia could fill. This would strengthen Australia's position to negotiate for continued access to frontier models and oversight of how they are developed.

Building inference capacity would reduce Australia's dependence on foreign compute and allow us to provide secure, trusted AI to regional partners.

Copyright is a barrier to both. Government must develop an Australian copyright approach that enables AI training and inference to occur in Australia while supporting creators. This does not need to be the perfect solution—rights holders can choose to opt out and negotiate their own agreements.

2. **The information and expertise to act on AI risks.** The Australian Government lacks the information and capability needed to protect Australians from AI risks.

Government should require frontier AI developers to be transparent about the risks of their models, and to report critical safety incidents.

Australia should make its AI Safety Institute world-leading by ensuring it is properly funded and making meaningful contributions to one or more critical areas of technical AI safety, such as compute verification.

3. **Regulations fit for the technology.** Regulators are not managing some emerging AI-specific risks.

Government should mandate gene synthesis screening to reduce AI-enabled biosecurity risk. The Department of Agriculture already has the power to act under the *Biosecurity Act 2015*, at near-zero cost.

Government should establish who is liable when AI agents cause harm.

If Australia does not act soon, we will inherit the decisions made offshore by very few countries, companies, and people. The window for securing our national interests in the age of AI is closing fast.

AI is more capable, and consequential, than people realise

In 2024, an AI model discovered a 'zero-day' for the first time—a previously unknown software flaw.

Eighteen months later, a new AI model discovered thousands of them across the world's major operating systems.

Two months after that, during a safety test, around a thousand AI agents broke out of containment, chained together zero-days and stolen credentials, reached the open internet, and hacked another company.

AI is moving fast

Many people are familiar with AI chatbots—often the free versions—but the frontier extends far beyond those systems. AI is approaching and surpassing human performance in a wide range of domains, including scientific reasoning, coding, cybersecurity, and persuasion.¹ These tools are being used by militaries to choose targets,² hospitals to help diagnose patients,³ laboratories to design new drugs,⁴ and companies to automate entry-level work.⁵ Frontier AI models are now needed to protect critical infrastructure from the cyber threats those same models pose when in the wrong hands.⁶

AI capabilities will continue improving. Progress has followed a consistent trend. More compute, more data, and better algorithms produce more capable models.⁷ Enormous capital is flowing into all three. This year, over US\$650 billion was invested in AI infrastructure.⁸ In 2025, Nvidia's revenue from AI chip sales alone was four times larger than Australia's defence budget.⁹ AI companies are offering up to tens of millions of dollars, per researcher, to draw in talent.¹⁰ The two largest economies have made AI a national priority. The United States is trying to 'win the race'¹¹ against China, and China is aiming to be the 'primary AI innovation center' by 2030.¹²

The pace of improvement will increase as AI is used to build itself. Anthropic's AI models wrote more than 80% of the company's code in 2026. This had risen from low single digits sixteen months earlier.¹³ Leading AI companies are explicitly working towards 'recursive self-improvement'—AI that autonomously designs and develops its successor.¹⁴ If they succeed, and each model builds a more capable successor, progress could accelerate sharply.

¹ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#). Department for Science, Innovation and Technology.

² Abraham, Y. (2024). ['Lavender': The AI machine directing Israel's bombing spree in Gaza](#), +972 Magazine.

³ US Food and Drug Administration. (2026). [Artificial intelligence-enabled medical devices](#).

⁴ Insilico Medicine et al. (2025). [A generative AI-discovered TNIK inhibitor for idiopathic pulmonary fibrosis: a randomized phase 2a trial](#). Nature.

⁵ Brynjolfsson, E., Chandar, B. & Chen, R. (2026). [Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence](#). Stanford Digital Economy Lab.

⁶ Anthropic. (2026). [Project Glasswing: Securing critical software for the AI era](#).

⁷ The length of tasks agents can complete unaided has been doubling every seven months: Kwa et al. (2025). [Measuring AI Ability to Complete Long Tasks](#). METR.

⁸ Reuters. (2026). [Big Tech to invest about \\$650 billion in AI in 2026, Bridgewater says](#).

⁹ Epoch AI. (2026). [Data on AI chip sales](#); SIPRI. (2026). [SIPRI Military Expenditure Database](#). Stockholm International Peace Research Institute.

¹⁰ Ho, A. (2026). [The economics of superstar AI researchers](#), Epoch AI.

¹¹ <https://www.ai.gov/action-plan>

¹² State Council of the People's Republic of China. (2017). [New Generation Artificial Intelligence Development Plan](#), (trans. Webster, G. et al), DigiChina.

¹³ Favaro, M. & Clark, J. (2026). [When AI builds itself](#). Anthropic Institute.

¹⁴ Favaro, M. & Clark, J. (2026). [When AI builds itself](#). Anthropic Institute; Altman, S. & Pachocki, J. (2026). [Built to Benefit Everyone: Our Plan](#). OpenAI.

Greater capability brings greater risk

Advanced AI can be deliberately misused, it can malfunction (including escaping human control), and it can cause systemic harms to society, such as labour-market disruption and the concentration of power.¹⁵

AI is already being misused

We're already seeing AI-driven harms, particularly in the cyber domain. In April this year, Anthropic's Mythos Preview model found thousands of high-severity flaws across the world's major operating systems and web browsers.¹⁶ The model was too dangerous to release publicly. Access was restricted to a small group of organisations. Models from other companies have similar cyber capabilities. OpenAI's latest Astra model has 'critical' cyber capabilities—it can find and exploit previously unknown security flaws across “many well-protected systems” without a person guiding each step.¹⁷ Such capabilities threaten Australian infrastructure like hospitals, water treatment plants, and the power grid.

AI cybersecurity capabilities are already being used against governments. In July 2026, suspected China-linked attackers ran a near-autonomous AI cyberattack on the Taiwanese government.¹⁸ The agents mapped 21 Taiwanese government systems, hacked into 85 user accounts, and extracted 2,500 personnel records. This attack reached a government email system, Taiwan's nuclear safety agency, and at least seven energy companies.¹⁹ Australia is vulnerable to similar AI-orchestrated attacks.

Biosecurity is another domain where AI could cause significant harm. All frontier developers warn their models could meaningfully help a non-expert create a biological weapon. AI systems can now provide expert-level guidance on biological weapons development,²⁰ and outperform PhD-level virologists on questions about highly technical laboratory procedures.²¹ AI-assisted pathogen design is also moving fast. Biological foundation models have generated genomes for viruses (bacteriophages) not found in nature, with some performing better than their natural counterparts.²²

As this submission was being written, Anthropic published five case studies of researchers using its models for dual-use biological work, including gain-of-function for the chikungunya virus and several weeks of study planning on mammalian adaptation in H5 avian influenza. In each case, researchers evaded Anthropic's access controls.²³

Australian experts are worried. In May 2026, nearly 150 AI experts, biosecurity experts, and members of the public called on the Australian Government to address AI-biosecurity risk.²⁴

AI is becoming harder to control

Models have begun to act in unintended ways, including breaking into systems they were not authorised to access. In July 2026, OpenAI disclosed that two of its models broke out of a secure testing environment, obtained internet access by exploiting a zero-day vulnerability in third-party software, and hacked into the company Hugging Face.

¹⁵ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#).

¹⁶ Anthropic. (2026). [Project Glasswing: Securing critical software for the AI era](#).

¹⁷ OpenAI. (2026). [GPT-6 Astra system card](#).

¹⁸ Blanchard, B. & Satter, R. (2026). [AI-driven hacking campaign targets Taiwan government agencies](#). Taipei Times.

¹⁹ Dream. (2026). [Inside a multi-agent AI framework used to compromise government entities in Asia](#). DREAM Lab Threat Research.

²⁰ Bengio, Y., et al. (2026). [International AI Safety Report 2026](#).

²¹ Götting, J., et al. (2025). [Virology Capabilities Test \(VCT\): A multimodal virology Q&A benchmark](#). arXiv.

²² King, S. H., et al. (2025). [Generative design of novel bacteriophages with genome language models](#). bioRxiv.

²³ Anthropic (2026). [Detecting and countering misuse of AI: September 2026](#). Anthropic.

²⁴ Australians for AI Safety. (2026). [Activate existing biosecurity powers to address AI-enabled risks](#).

Over a thousand individual AI agents collaborated on a message board in an attempt to cheat on the evaluation they were being tested on.²⁵ The models took roughly 17,600 actions over four days to hack Hugging Face and steal the answers to the test.²⁶ Anthropic, Meta, and the UK AI Security Institute have separately reported that their models, during testing, hacked computer systems.²⁷ We only know what has been reported voluntarily. There may be similar incidents happening right now.

We also see AI agents manipulating humans. In July 2026, the UK AI Security Institute discovered that an AI agent attempted to insert malicious code into a public open-source project. To get the code approved, the AI researched the project's human maintainers, created multiple fake online identities, and used them to pressure a person into adopting the malicious code. When the code was challenged, the agent edited its earlier activity to appear harmless and considered adopting a fresh identity to continue its deception.²⁸

The harms we see today are warning shots

More capable models will be trusted with more consequential decisions. They will be woven tighter into the economy and critical infrastructure. When they fail or are misused, the damage will be greater and harder to reverse. Maintaining oversight will become harder, as models act more autonomously and with reasoning humans cannot fully understand.²⁹ AI harms are rapidly intensifying and becoming systemic as AI reshapes our economy, national security, and distribution of power.³⁰ Warnings are being issued from AI companies³¹ and scientists³² to financial regulators³³ and the heads of Five Eyes Security agencies.³⁴

Australia is exposed to these risks, and unprepared for them. To manage these risks and protect national interests, Australia must have:

- 1) influence over how AI develops,
- 2) the information and expertise to act on AI risks, and
- 3) regulations fit for the technology

Recommendation 1: The Committee's report should acknowledge that AI poses catastrophic and potentially existential risks, and that a proportionate government response is warranted.

AI capability is advancing rapidly, and greater capability brings greater risks—from cyberattacks on critical infrastructure to biological weapons development and the loss of human control. The Australian Government should explicitly acknowledge these risks as part of its response.

²⁵ METR & Redwood Research. (2026). [Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident](#).

²⁶ Larcher, H. et al. (2026). [Anatomy of a frontier lab agent intrusion: A technical timeline of the July 2026 incident](#). Hugging Face; OpenAI. (2026). [Hugging Face model evaluation security incident](#).

²⁷ Anthropic. (2026). [Investigating three real-world incidents in our cybersecurity evaluations](#); AI Security Institute. (2026). [Unsanctioned agent behaviour during cyber evaluations](#).

²⁸ AI Security Institute (2026). [Security Incident INC-2026-07-28-01](#). UK Department for Science, Innovation and Technology.

²⁹ OpenAI reported their ability to monitor model's reasoning (their 'chain of thought') has decreased. OpenAI. (2026). [GPT-6 Astra system card](#).

³⁰ Slattery, P., Saeri, A. K., Graham, J., Noetel, M., & Thompson, N. (2026). [Prioritization of Risks from Artificial Intelligence: A Delphi Study of 272 International Experts](#). MIT FutureTech and the University of Queensland.

³¹ Employees of frontier AI companies. (2026). [Pacing the frontier](#).

³² Bengio, Y., et al. (2026). [International AI Safety Report 2026](#).

³³ Australian Securities and Investments Commission (2026). [ASIC and APRA warn frontier AI awareness must turn to action](#).

³⁴ Australian Signals Directorate. (2026). [Five Eyes cyber security agencies statement](#).

Australia needs influence in the age of AI

To manage these risks, Australia needs influence over how AI develops. That means finding valuable positions in a world with highly capable AI, and using those positions to secure our national interests.

A position is valuable if it secures Australia a place in the economy and supply chains that will underpin future prosperity, gives us something to offer our allies and regional partners that they cannot easily obtain elsewhere, or reduces our dependence on decisions made offshore.

Good Ancestors has identified ten positions Australia could pursue, and assessed the value and feasibility of each (see below).³⁵ Value is how much Australia gains if we attain the position, while feasibility is how likely Australia is to attain it. This section outlines two: hosting frontier AI training (**Position 1**) and building inference capacity (**Position 2**). We found that building a world-leading AI Safety Institute (**Position 7**) was the most tractable action for Australia (see page 13).

Ten valuable positions in the age of AI

Good Ancestors identified and analysed ten ways Australia could establish its place in an AI future:

1. Attract frontier AI training: Australia could attract one or more frontier AI companies to make their future AI models in Australia.
2. Build inference capacity for Australia and the region: Australia could have enough domestic AI inference compute to meet domestic demand and be a net exporter.
3. Secure industries across the AI value chain: Australia could secure multiple valuable places across the AI supply chain, beyond hosting AI training and inference compute.
4. Fund narrow AI models in areas of national strength: Australia could support the development of truly sovereign small or narrow AI models in areas of national strength.
5. Develop a sovereign fast chaser: Australia could support the development of an Australian "fast-chaser" AI model that aims to have most of the performance of frontier models at a fraction of the cost.
6. Coordinate with middle powers on AI: Australia could use its influence and trusted relationships to coordinate middle powers on reducing AI risks and mitigating power concentration.
7. Fund a world-leading AI Safety Institute: Australia could have a world-leading AISI that makes a meaningful contribution in one or more critical areas of technical AI safety, in ways legible to our allies and frontier AI companies.
8. Drive AI adoption across industry: Australia could accelerate the safe and responsible adoption of AI across domestic industry.
9. Invest in AI-proof industries: Australia could invest in industries unrelated to AI but that are likely to become much more valuable in a world with cheap intelligence available at scale.
10. Host the new industries that AI will create: Australia could prepare for a large expansion in energy, land use, and industrial capacity so that the new industries AI creates are built here.

³⁵ Good Ancestors. (2026). [Australia's position in the age of AI](#).

The below matrix shows how each position scored on value and feasibility.



Hosting frontier AI training gives Australia more of a say

Hosting frontier AI training may be the most valuable action Australia could take to establish its place in an AI future.³⁶ Creating frontier models, rather than just running them, is a critical place in the AI supply chain. Only a handful of countries will host training, as it requires large, concentrated amounts of compute and power that few countries can support, and does not need to be globally distributed. Australia, with our land, renewables, and stable governance, has attracted interest as a frontier training destination.³⁷

Hosting training would put Australia in a stronger position to negotiate for terms in the national interest, including securing access to frontier models, overseeing how they are developed, and attracting talent.

Securing access to frontier models

Australia has no domestic frontier models and has already struggled to gain access to the most advanced models.³⁸ Even the UK's AI Security Institute did not secure pre-release access to Mythos 5.1.³⁹ Frontier access will only get harder to rely on as demand for compute to train and run frontier models grows, security concerns increase, and access becomes a strategic asset that Australia has little basis to bargain for.

Losing access to frontier AI has serious implications for national security, the economy, and public-interest research.

³⁶ Good Ancestors. (2026). [Australia's position in the age of AI](#).

³⁷ Wilson, C. (2026). [AI giant Anthropic wanted 'massive' data centre capacity in NSW, emails reveal](#). ABC News.

³⁸ Anthropic's Mythos cybersecurity model was initially only released to a restricted group of mostly US-based organisations. This did not include the Australian Government. When a public version was released (Table 5), it was available for only three days before a US Government export directive required Anthropic to suspend access for foreign nationals.

³⁹ Kelly, R. (2026). [Anthropic reportedly withholds access to Mythos 5.1 from UK safety testing body](#). IT Pro.

For national security, effective defence will require AI at least as capable as that available to Australia's adversaries.⁴⁰ An adversary using more capable AI, including for cyber operations, intelligence, or autonomous systems, is at a significant advantage. Without access, Australia cannot harden systems of national significance, from health systems to the power grid, before adversaries use the same capabilities against us. As AI cyber capability continues to accelerate, countries with access to frontier AI cyber capabilities will be able to dominate those without access.

For the economy, businesses with access to more capable AI models will outcompete those only using less capable AI models. As AI drives more output, the economic gap between AI “haves” and “have-nots” will grow. If Australia wants to create valuable systems and products built on AI models, as the National AI Plan envisages, we need Australians to have access to the most capable AI models where that capability matters the most.

For government and research, the Australian AI Safety Institute needs frontier-model access to evaluate capabilities and risks. Australian universities and research institutions also need access to do high-quality research and build local expertise. Scientific breakthroughs are likely to come from frontier models, and be made by research groups with that access.⁴¹ Conversely, researchers with limited compute and less capable models will fall further behind.

Hosting training strengthens Australia's ability to negotiate continued access. Companies that have invested billions in training clusters here, and that rely on Australian infrastructure and regulatory cooperation, have reason to support Australia's continued access to the frontier. The US government would also have more reason to back Australian access, as an ally hosting important AI infrastructure the US has a stake in.

Australia should seek access at three levels:

- 1) market access for the Australian public, equivalent to the US,
- 2) access to restricted or sensitive models for critical infrastructure operators and other defenders, and
- 3) Government access, equivalent to the US government for national security purposes.

Securing oversight of AI development

Hosting training could give Australia oversight of how frontier models are developed. Australia could ask AI companies training here to disclose the capabilities and risks of the models they train, report critical safety incidents, and let the Australian AI Safety Institute (AISI) test models before deployment.

Australia could try to impose some of these requirements through regulation, even if training does not happen here—as New York and the EU do.⁴² But as a smaller market, companies may be more likely to comply if they are heavily invested here and rely on continued regulatory cooperation.

Further, some obligations only attach to infrastructure physically located here, like the security of the facility itself and oversight of training as it happens. The *Security of Critical Infrastructure Act 2018* is one example. The Act requires critical infrastructure assets located in Australia, including data storage and processing assets, to manage risks and report cyber incidents.⁴³ The Department of Home Affairs is currently reviewing amendments that would ensure AI data centres are covered by the Act.⁴⁴

⁴⁰ Phillips-Robins, A., Tawil, T. & Winter-Levy, S. (2026). [The compute coalition: How to build the future of AI in the free world](#). Carnegie Endowment for International Peace.

⁴¹ OpenAI (2026). [On the Navier–Stokes Millennium Prize Problem](#).

⁴² EU AI Act (Regulation (EU) 2024/1689), article 55; New York State Legislature. (2025). *Responsible AI Safety and Education (RAISE) Act* (S6953-B).

⁴³ Parliament of Australia. (2018). *Security of Critical Infrastructure Act 2018*.

⁴⁴ Department of Home Affairs. (2026). [Proposed amendments to streamline and modernise the SOCI Act 2018](#).

Australia could go further and independently verify what data centres on our soil are doing—including whether a facility is training or running models, and whether it is meeting agreed conditions. This is known as ‘compute verification’, and is a promising area where the AISI could develop expertise.

Recommendation 2: Attract AI training to Australia and apply national interest conditions.

Australia should adopt the policy settings necessary to attract frontier AI training compute to Australia. Leveraging our national strengths, like abundant land and energy, security, and stability, Australia could secure commitments including:

- Ongoing access to frontier AI models including public/market access equivalent to the US, access to restricted or sensitive models for critical infrastructure operators and other defenders, and Australian Government access for national security purposes.
- Oversight over AI development in Australia, including pre-deployment testing of models trained in Australia by the Australian AI Safety Institute, disclosure of model capabilities and risks, and reporting of critical safety incidents to Government.
- A portion of compute allocated for research and innovation on favourable terms (e.g., to Australian researchers, not-for-profits, and start-ups).

Other potential terms could include requiring some AI engineers and researchers to be based in Australia, ensuring AI data centres have high security standards, notifying Government of large training runs conducted in Australia, and ensuring fair contributions are made via taxation and other mechanisms.

Building inference capacity would boost our economy and security

While building AI models sits further up the value chain, running AI models (‘inference’) and offering them regionally supports Australia’s economy and security.

As AI is adopted across the economy, including in critical infrastructure, failure to meet domestic inference demands introduces a foreign dependency and exposes Australia to additional AI security risks. Countries that control our inference compute could drive up costs or remove access. They could also exploit their access to create vulnerabilities, including by substituting the model for a backdoored variant, or extracting and leaking prompts, outputs, and model weights.⁴⁵

Australia has a compute deficit. Our current best estimate is that **Australia’s inference compute supply meets less than half of domestic demand**, and we only host about 0.4% of global compute supply.⁴⁶ Even accounting for newly announced AI data centres, we expect **Australia will meet less than a quarter of its domestic inference demand by the end of 2027**. This is due to an incoming ‘compute crunch’, with global demand growing at 10x per year, while supply only increases by 3x.⁴⁷

Australia should be able to meet domestic inference demand, including as that demand grows, with a surplus. That surplus would let Australia provide secure, trusted AI to regional partners for whom hosting compute at national scale is not viable. Our neighbours in the Pacific and parts of Southeast Asia may lack the power generation, grid capacity, land, or capital to build sovereign AI infrastructure. However, their governments, banks, hospitals, and utilities will still want to adopt AI.

⁴⁵ Cai, W. et al. (2025). [Are you getting what you pay for? Auditing model substitution in LLM APIs](#). arXiv:2504.04715

⁴⁶ Good Ancestors. (2026). [Australia’s position in the age of AI](#).

⁴⁷ Emberson, L. & Sevilla, J. (2026). [Is a compute crunch coming?](#). Epoch AI.

Meeting that demand would reduce their dependence on alternatives, such as Chinese AI infrastructure, and give Australia standing in how AI is governed across the region.⁴⁸

Pursuing this position requires new work to measure Australian supply and demand and set a policy target to ensure we are a net compute exporter.

Recommendation 3: Ensure Australia can meet its own inference demand and serve the region.

Australia currently meets less than half of its domestic inference demand, and is on track to meet less than a quarter of demand by the end of 2027. The Australian Government should:

- Measure Australian inference compute supply and demand on an ongoing basis.
- Set a target for Australia to become a net exporter of inference compute, and use the target to guide decisions on energy supply, grid connection timelines, and planning approvals.

Capturing the benefits of AI training and inference requires a copyright solution

Hosting AI training and inference in Australia might be the best way for Australia to secure a valuable position in the AI value chain and protect itself from the broader risks of AI. Copyright is a barrier to both.⁴⁹

Training frontier AI models in Australia is not viable under our copyright law

Australia has no fair use or text-and-data-mining exemption permitting training on publicly available material, unlike the United States, the European Union, Singapore, Israel, and Japan.⁵⁰ Under Australian copyright law, AI companies and rights holders must strike licensing deals covering all content used to train a model.

AI companies don't see licensing as feasible, so they generally have not entered into negotiations. Collective licensing and reciprocal agreements only cover some categories of material and do not reach a large amount of content—blogs, forums, and computer code—that developers want to train on. AI companies do not think it is worth licensing some content if they don't have a viable pathway for other content. Australian copyright also protects all works, not only Australian ones, so training here would require licensing solutions covering the entire global corpus.

Rights holders want control over whether their works are used to train AI, and compensation for when they are. Rights holders are reluctant to accept any alternative beyond licensing. Rights holders have also expressed distrust and animosity towards AI companies.

This leaves Australia's national interest dependent on a negotiation between two third parties that is unlikely to progress. For that reason, the Australian Government should intervene to facilitate a deal that allows AI training while supporting Australia's national interests, including the interests of Australian creators.

⁴⁸ For discussion of this concern, see, e.g., Forbes, K. & Maughan, C. (2026). [ASEAN must strike a balanced AI relationship with China](#).

⁴⁹ Training and inference are run in data centres, which carry other important considerations like energy demand, water use, and community impacts. We address these considerations in the Senate inquiry into AI and data centres: Good Ancestors. (2026). [Hosting AI infrastructure offers Australia more than economic benefits](#).

⁵⁰ United States: Copyright Act 1976, 17 U.S.C. §107; European Union: Directive (EU) 2019/790, art. 4; Singapore: Copyright Act 2021, s. 244; Japan: Copyright Act 1970, art. 30-4; Israel: Copyright Act 2007, s. 19, 22, 26, and the Ministry of Justice. (2022). [Opinion: Uses of Copyrighted Materials for Machine Learning](#).

YouGov Polling shows a majority of Australians back some sort of copyright reform that enables frontier AI training.⁵¹ In April 2026 polling, 61% of Australians said they would prefer to enable AI training under a different copyright arrangement that still supports creators, against 15% who would keep current laws even if it meant AI companies trained overseas.

Good Ancestors has consulted on one approach (Attachment A),⁵² and we provide a summary of initial consultation outcomes (Attachment B). Other uniquely Australian approaches may be possible.

Copyright also puts inference on unclear legal footing

Independent of AI training, copyright reform may be necessary to enable current AI use. Australian AI data centres mostly do inference computing. During inference, copyrighted material gets copied. This can include the AI model accessing works, like news or reports, and storing them in memory. Most assume this does not infringe copyright because section 43B of the Copyright Act allows a 'temporary' reproduction of a copyright work where the copy is a 'necessary' and 'incidental' part of a technical process.

However, making copies during inference may not be temporary, necessary, or incidental. Models potentially hold copyrighted material in 'context windows' for extended periods. 'Cross-chat memory' lets models recall material from previous conversations. Longer 'chains of thought' allow models to reason better about copied material. These design choices are made to improve performance, and are perhaps not a necessary or incidental part of AI technology.⁵³ Possible technological developments, like 'continual learning' where a model's weights shift in response to user behaviour, could further blur the boundary between use and training.

The consequences of inference being found to infringe copyright in Australia would be severe. AI use and data centre investment rest on the assumption that the day-to-day use of AI does not infringe copyright. If that assumption were overturned, adoption would stall and ongoing investment would be undercut.

⁵¹ Good Ancestors. (2026). [AI Training and Copyright: YouGov polling commissioned by Good Ancestors](#).

⁵² Good Ancestors. (2026). [Building AI while Compensating Creators: Consultation Paper on an Australian AI Training Scheme](#).

⁵³ Wei, J. et al. (2022). [Chain-of-thought prompting elicits reasoning in large language models](#); Denain, J.-S. & Barry, A. (2026). [Have AI capabilities accelerated?](#) Epoch AI; Epoch AI. (2026). [Trends in artificial intelligence](#).

Copyright and AI have a workable solution

Australian copyright law requires licences between AI companies and rights holders before training can happen. But rights holders and AI companies don't see eye-to-eye and aren't on track to make deals at the scale needed to train AI in Australia.

Under this status quo, the Prime Minister's goal of building AI in Australia and securing multiple places in the AI supply chain will not be achieved.

Australia doesn't need one perfect solution. Instead, we can find a new, optional, pathway and let rights holders choose between the new pathway and the status quo.

The Committee should recommend:

- An optional pathway that makes training feasible by not requiring licensing for every global work.
 - We proposed a fund-and-permit model (see Attachment A).
- That companies training AI in Australia allow Australian rights holders to opt out.
 - Opt-out should be practical, including using the best available technology, such as robots.txt and content ID systems.

This gives Australian rights holders a choice:

- If the new pathway works for them, they can follow it.
- If they don't like the new pathway, they can opt-out and either enter commercial negotiations under existing laws or not be involved in AI training.

The approach does not weaken copyright law because it is *additional* to the status quo. Anyone who opts out retains existing copyright protections and can negotiate their own agreements. Australian creators and rights holders have control—they choose which path to take—and they gain a new way to be paid if Government pursues a fund-and-permit model.

Recommendation 4: Develop copyright pathways that allow AI training and inference to occur in Australia while supporting creators.

The benefits of hosting AI training depend on AI training being viable here. Existing licensing approaches are not on track to enable AI training.

- Government should progress an Australian approach—not a text-and-data-mining exception—that supports rights holders while allowing training to occur in Australia.
- Good Ancestors' consultation paper sets out one potential model, but other creative and uniquely Australian approaches may also be possible.

Recommendation 5: Clarify the position of inference under Australian law and ensure copyright law is fit for purpose.

The Government should seek legal advice on whether AI inference falls within section 43B of the *Copyright Act 1968*, and should clarify its treatment in any copyright reform rather than leaving the question to litigation.

Australia needs information and expertise to act on AI risks

Australia not only needs influence over how AI develops, but the capacity to see and understand risks.

Australia needs visibility into AI risks

Australia is flying blind. We do not have information about the most capable AI models. We need AI companies to be transparent about their models, and to report critical incidents to policymakers.

Transparency means requiring developers of the most powerful models to disclose what their models can do and the risks they carry. In practice, this means a technical description of the model and its capabilities, an assessment of its dangerous capabilities and what the company has done to address them, and a governance framework describing how the developer is meeting safety standards. This is currently voluntary in Australia, and we benefit incidentally from what California, New York, and the EU require. Our national interest can't rest on the goodwill of AI companies. These companies produce much of this information already, and Australia should require it—as other jurisdictions do—so we can see the risks we're trying to manage.

Transparency is about what a model *could* do. Incident reporting is about the negative impacts a model actually has. **Incident reporting is a duty to notify Government when a model causes a critical safety incident**, like helping build a biological weapon, autonomously conducting a cyberattack, or disrupting critical infrastructure. Australia needs legislation to define what counts as a critical safety incident, set a notification period, and require companies to report this information to the Australian Government. Australia operates analogous systems for aircraft, medicines, and consumer products. No Australian regulator currently captures this information for underlying AI models, even though it's required in other jurisdictions like California⁵⁴ and the EU.⁵⁵

Australia needs expertise to understand and respond to risks

Australia needs the sovereign capability to evaluate, advise on, and mitigate frontier AI risks. That capability is mainly housed in the Australian AI Safety Institute (AISI).

Australia should make our AISI world-leading—not only because that improves our ability to manage AI risks, but also because a valuable contribution to global AI safety research strengthens our alliances and our position to shape AI development. The UK AISI demonstrates this. Because of their technical expertise, they secured agreements with OpenAI and Anthropic to evaluate models before deployment. These collaborations allow them to understand the state of frontier technology, advance innovation and security in their AI industry, and prompt improvements in the most widely-used AI systems.⁵⁶

Australia's AISI should develop world-leading capability in one or more areas of AI risk. These areas should be useful to global research, our allies, and frontier AI companies. A valuable contribution gives others reason to share information with Australia in areas of expertise. As noted earlier in 'Securing oversight of AI development', compute verification is a strong candidate for specialisation. Expertise would let Government independently confirm what data centres on Australian soil are training or running, AI companies are calling for these verification tools,⁵⁷ and in the longer term, verification will be needed for any international agreement on how AI is developed.

⁵⁴ California State Legislature. (2025). [Senate Bill 53: Artificial Intelligence Models: Large Developers](#).

⁵⁵ European Parliament & Council of the European Union. (2024). [Regulation \(EU\) 2024/1689 \(AI Act\), Article 73](#).

⁵⁶ AI Security Institute. (2025). [How we're working with frontier AI developers to improve model security](#).

⁵⁷ Anthropic. (2026). [When AI builds itself](#); OpenAI. (2026). [Built to benefit everyone: our plan](#).

To be a leader, Australia's AISI needs sufficient funding. The AISI receives an average of \$7.5 million a year, with funding set to decline over time—contrary to Government's view that the impacts of AI will increase over time. With the current level of resourcing, the AISI cannot properly fulfil its mandate. It will lack the technical staff needed to conduct evaluations of frontier AI systems, the capacity to stay ahead of advancing capabilities, and the personnel to meaningfully support dozens of regulators managing AI risks across healthcare, finance, transport, critical infrastructure, and other sectors. To satisfy its mandate, **AISI funding should increase to at least \$50 million per year plus compute access, roughly matching the UK's world-leading AISI when adjusting for population or GDP.**⁵⁸

Recommendation 6: Get Australian policymakers the information they need.

The Australian Government does not have enough information to appropriately manage frontier AI risks. The Australian Government should:

- Create transparency obligations that require developers of the most capable models to disclose what their models can do and their risks.
- Create incident reporting obligations that require developers to report critical safety incidents to the Australian Government.

Recommendation 7: Make the Australian AI Safety Institute world-leader.

The Australian Government should:

- Increase the Australian AI Safety Institute funding to at least \$50 million a year.
- Direct the Institute to build world-leading capability in areas that give allies and Australian regulators the information they need, such as compute verification.

Australia needs regulations fit for AI

Australia's regulations must reflect AI's speed and consequences. The National AI Plan signalled the Government wasn't pursuing a whole-of-economy AI Act, and was instead asking regulators to manage AI risks in their domains. The Plan says:⁵⁹

The government is actively monitoring emerging risks. Where necessary, we will take decisive action to ensure safety and accountability as new technologies and frontier AI systems emerge. Existing regulators will continue to identify and manage harms and report any gaps in laws to the AISI. We will respond to emerging risks including bias, privacy breaches, disinformation and cyber threats. If more regulation is needed to address bad actors or broader harms, the government will not hesitate to intervene.

This approach is under pressure. There are specific examples where existing regulators are not managing emerging AI risks, and the government's pledge to 'intervene without hesitation' should be activated.

⁵⁸ Good Ancestors (2026). [Pre-Budget Submission 2026–27](#).

⁵⁹ Department of Industry, Science and Resources. (2025). [National AI Plan: Keep Australians safe](#). Australian Government.

Mandate gene synthesis screening to address AI-bio risk

AI is eroding the barriers to building a biological weapon. Models can already troubleshoot technical laboratory protocols,⁶⁰ design working viruses not found in nature,⁶¹ and meaningfully help non-experts build biological weapons.⁶² AI companies are now reporting instances of actors evading guardrails to do dual-use biology research.⁶³

Biosecurity risks from AI are not currently catastrophic, but will escalate quickly. In cyber, less than two years separated the first warning signs from AI models hacking real companies. Australians for AI Safety first raised AI-biosecurity concerns with the Australian government on 10 May 2024—more than two years ago.⁶⁴ Life science researchers, AI companies, and national security experts are warning that urgent action is needed now.⁶⁵

Fortunately, Australia already has the power to act. While AI may soon be able to help someone build a bioweapon, they still need the genetic material, called ‘synthetic nucleic acids’. In Australia, given a lack of domestic industry, this material must be imported. Under the *Biosecurity Act 2015*, the Director of Biosecurity could require that synthetic nucleic acids come only from providers that screen orders for dangerous sequences and verify customers. No new legislation is needed, and major synthesis providers, including Twist Bioscience, IDT, and GenScript, already screen voluntarily. The US,⁶⁶ New Zealand,⁶⁷ the EU,⁶⁸ and the UK⁶⁹ are all moving on screening. In May 2026, almost 150 signatories asked the Minister for Agriculture to act.⁷⁰ As of September 2026, no action has been taken.

The consequences of acting slowly on biosecurity risks are far worse than those of acting slowly on cybersecurity risks. A cyber ‘warning shot’ means a company gets hacked. An engineered pathogen ‘warning shot’ could mean a global pandemic. Australia has committed to reducing biosecurity risk to a very low level, and the current science is clear that the risk is far higher than ‘very low’.⁷¹ The National AI Plan asks regulators to identify and close gaps in their own domains. This gap is clear in the scientific literature, has been identified for the regulator, the power exists, the cost of action is near zero, but action proportionate to the risk has not been taken.

⁶⁰ Götting, J., et al. (2025). [Virology Capabilities Test \(VCT\): A multimodal virology Q&A benchmark](#).

⁶¹ King, S. H., et al. (2026). [Generative design of bacteriophages with genome language models](#). *Science*, 393(6811).

⁶² Bengio, Y., et al. (2026). [International AI Safety Report 2026](#).

⁶³ Anthropic. (2026). [Detecting and countering misuse of AI: September 2026](#).

⁶⁴ Australians for AI Safety. (2024). [Senate Inquiry on Adopting Artificial Intelligence](#). Letter to the Senate Select Committee on Adopting Artificial Intelligence, 10 May.

⁶⁵ ScreenDNA (2026). [An Open Letter in Support of Mandatory Nucleic Acid Synthesis Screening and Recordkeeping](#).

⁶⁶ The White House. (2025). [Improving the safety and security of biological research](#).

⁶⁷ New Zealand Parliamentary Counsel Office. (2025). [Gene Technology Bill 110-2](#) (2024).

⁶⁸ European Commission. (2025). [Proposal for a regulation to establish measures to strengthen the Union's biotechnology and biomanufacturing sectors](#).

⁶⁹ UK Department for Science, Innovation and Technology. (2024). [UK screening guidance on synthetic nucleic acids for users and providers](#).

⁷⁰ Australians for AI Safety. (2026). [Activate existing biosecurity powers to address AI-enabled risks](#).

⁷¹ *Biosecurity Act 2015* (Cth), s 5.


Recommendation 8: Require imported synthetic nucleic acids to come from providers that screen.

The Director of Biosecurity can impose the screening requirements needed to reduce catastrophic biosecurity risk now, under the Biosecurity Act 2015. No new legislation is needed and the cost is very low. The Australian Government should:

- Add a condition to import permits in the *Biosecurity Import Conditions system* (BICON) requiring that synthetic nucleic acids come only from providers that screen orders for dangerous sequences.

Establish who is liable when AI agents cause harm

In August 2025, Good Ancestors published an *Australian AI Legislation Stress Test*.⁷² The Stress Test asked 64 experts in AI, law, cybersecurity, national security, and policy to assess five AI threats against current Australian law. Two of the five related to AI agents. Experts said it was ‘probable’ that AI agents would cause deaths or tens of millions of dollars in economic damage annually within the next 5 years. Four in five experts rated current Government measures for agent threats as inadequate. The legal analysis found that no single regulator is tasked with addressing them.



Government itself has identified regulatory gaps with AI agents, particularly around liability. Treasury's Review of AI and the Australian Consumer Law, published in October 2025, found that agentic AI requires further consideration of liability mechanisms, and that this consideration should occur in a coordinated way across legislative frameworks to avoid regulatory dislocation.⁷³ This has not occurred. Good Ancestors' submission to that review argued that consumer law should ensure AI developers are responsible for risks caused by their systems, rather than shifting risks to the Australian businesses and consumers who are not well placed to address them.⁷⁴

⁷² Good Ancestors, (2025). [Australian AI Legislation Stress Test: Expert Survey](#).

⁷³ Department of the Treasury. (2025). [Review of AI and the Australian Consumer Law: Final report](#). Australian Government.

⁷⁴ Good Ancestors. (2024). [Submission to the Review of AI and the Australian Consumer Law](#).

Other jurisdictions have taken action to address these risks. California's Transparency in Frontier Artificial Intelligence Act states that a frontier model carrying out a cyberattack without meaningful human oversight, or evading the control of its developer or user, is a catastrophic risk, and requires developers to report loss-of-control incidents to the state within 15 days.⁷⁵ New York's RAISE Act and Article 55 of the EU AI Act impose comparable reporting duties.⁷⁶

Since the Stress Test was published, both types of agent threat have occurred. In August 2026, a Melbourne man asked his AI agent to book a gym class.⁷⁷ The agent found a flaw in the gym's booking system, cancelled the reservation of the person ahead of him on the waitlist, and secured him a place. Although the harm was trivial, it shows the risk is real and the law is uncertain. Australia's computer offences turn on a person's mental state.⁷⁸ In this case, the user did not intend the harm, and the AI developer was unaware of the incident. ASD's advice appeared to blame the victim organisation.⁷⁹ The Hugging Face incident is a further example of this risk occurring and causing harm.

The National AI Plan committed Government to taking 'decisive action' to ensure accountability as risks emerge. This is not occurring.

Recommendation 9: Ensure accountability for harms caused by AI agents.

Existing regulations and legislative frameworks are inadequate. The Australian Government should:

- Direct regulators to clarify how their regulations apply to AI agents.
- Amend consumer law so that AI developers are responsible for what their models can do, including the risks their systems create. Liability should not shift to Australian businesses and consumers when they are not well placed to address those risks.

Conclusion

AI development is vastly outpacing Australian action. In cyber, less than two years separated the first AI-discovered vulnerability from swarms of AI agents breaking containment and hacking companies. Biosecurity risks are on the same trajectory. AI is reorganising the global economy and upending geopolitics.

Australia is exposed to all these risks and has a rapidly diminishing ability to address them. We do not build frontier models. We host only 0.4% of global compute. Our access to the most capable systems has already been granted and withdrawn at the whim of a foreign government. Our sovereign AI industry is dramatically underscaled. Our talent is leaving for the US and China. Our current best positions are at either extreme of the value chain—exporting critical minerals and importing AI models.

Australia still has time to secure valuable positions in the AI value chain. Our submission sets out some options. Other submissions will have other good ideas. We need to overcome the comparably minor hurdles and seize as many of these valuable positions as possible. Once we have national wherewithal, we must use it to protect and advance our national interests—while we still have national autonomy.

⁷⁵ California State Legislature. (2025). [Senate Bill 53: Transparency in Frontier Artificial Intelligence Act](#) (ss 22757.11–22757.13).

⁷⁶ New York State Legislature. (2025). [Responsible AI Safety and Education \(RAISE\) Act \(S6953-B\)](#); European Parliament & Council of the European Union. (2024). [Regulation \(EU\) 2024/1689 \(AI Act\)](#) (Article 55).

⁷⁷ Wilson, C. & Hobbins, R. (2026). [AI assistant hacks gym website in first known Australian autonomous cyber attack](#). ABC.

⁷⁸ Criminal Code Act 1995 (Cth), Part 10.7.

⁷⁹ Australian Signals Directorate. (2026). [When AI agents take unexpected actions](#). Australian Cyber Security Centre.

Attachment A: Building AI while Compensating Creators

Consultation Paper on an Australian AI Training Scheme

Executive Summary

Australia faces a decision point on AI training. Australia is an attractive destination for training frontier AI models — we have abundant renewable energy, stable governance, and a skilled workforce. But copyright is a barrier. Unlike some other countries, Australia has no tailored legal regime that makes frontier AI training viable here.

The status quo serves no one. Australian creators receive little or no compensation when their work is used to train AI overseas. AI investment is deterred. Australia loses both the economic benefits and the strategic advantage of hosting AI training infrastructure. Australia's ability to pursue its national interest depends on negotiations between third parties—global AI companies and global rights holders.

The Government has ruled out a “text and data mining (TDM) exception” but has signalled it is open to other approaches. The Copyright and AI Reference Group is considering options, including how to encourage fair, legal avenues for the use of copyright material in AI training.

This paper proposes an Australian AI Training Scheme. The scheme has two components:

An **AI Training Permit** would be available for AI training in Australia. Obligations would include paying a tiered fee into consolidated revenue, granting practical consent and control to rights holders via an opt-out mechanism, complying with the National Data Centre Expectations, and meeting other national-interest conditions. Permit holders would be authorised to train AI with the certainty that they would not be subject to copyright restrictions in respect of that training. Permit holders would also receive support in navigating federal, state and territory requirements to accelerate “time to power” for new data centres.

An **Australian AI Public Benefit Fund** would be established separately by Government, funded from consolidated revenue, to support Australian creators. Its design would draw on the *Public Lending Right Act 1985* — compensation for income loss, and support for Australian culture.

The permit and the fund operate separately. This avoids characterising the fund as a de facto licensing scheme, preserves the Government's ability to adjust each independently, and gives ongoing policy levers as global circumstances evolve.

This paper invites comment from creators, collecting societies and rights management organisations, AI companies, and the broader public (see consultation questions on p.9). Responses are sought within three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

Background

AI is becoming increasingly important to the global economy. Investment in AI infrastructure is globally significant – arguably the largest coordinated investment in modern history. Leading US technology companies alone are spending over US\$650 billion on AI infrastructure in 2026.

AI training is a distinct and critical element of the AI value chain. While many countries and regions will host AI "inference computing" – the hardware that runs AI models and serves local, time-sensitive demand – AI training compute is likely to be far more centralised. Training frontier AI models requires concentrations of specialised hardware and energy. Most countries will not have AI training compute. Where training happens matters, because the countries that host it capture greater economic value, strategic leverage, and a practical ability to shape how AI is developed.

Like any large infrastructure project, data centre construction and operation has challenges. These relate primarily to power consumption, but also land use, water, air quality, noise, and other local impacts. The National Data Centre Expectations set out how these issues should be addressed. AI companies are on the record, in Australia and overseas, agreeing to requirements of this kind, including covering grid infrastructure costs and offsetting consumer electricity price increases.

Copyright is the key outstanding barrier. Current Australian copyright law makes it uneconomic, and arguably logistically impossible, for AI companies to train AI models in Australia. Unlike the United States, Israel, the European Union, Singapore and Japan, Australia does not have a fair use or TDM exemption that would permit AI companies to train on publicly available material. Under the status quo, AI companies will continue to train in jurisdictions where training data is available at low or no cost.

Australian creators argue that AI companies should license their materials. AI companies are unlikely to follow this path for three reasons. First, the AI companies are not required to license the material in other jurisdictions. While there is legal uncertainty in some jurisdictions like the US, the law is unequivocal in other jurisdictions, like Japan, Singapore and Israel. Second, Australian copyright law protects all copyrighted material, not just Australian material. Training AI in Australia would require licensing copyrighted works globally – something AI companies believe is unworkable. Third, formally licensing publicly available copyrighted material could undermine favourable court rulings in other jurisdictions, particularly the United States, where AI training is currently treated as fair use.

The status quo serves no one well. Creators receive little to no payment and may not be able to withdraw permission when their material is used to train AI in the US or China, where most AI training occurs today. AI investment is deterred because current settings make training here uneconomic. And Australia's geostrategic position is weakened – without a place in the AI training value chain, Australia has limited leverage to shape the future of AI in the national interest.

Political context

The Australian Government has ruled out a TDM exception to copyright. In October 2025, Attorney-General Michelle Rowland stated:

While the Government is not considering a Text and Data Mining Exception, work is underway to ensure that Australia is prepared for future copyright challenges emerging from AI – so that Australian creators are protected and supported while unlocking new uses of copyright material.

The Government has since tasked the Copyright and AI Reference Group with work across three priority areas, including to "encourage fair, legal avenues for using copyright material in AI." The Attorney-General has framed the challenge as one of balance:

Artificial Intelligence is an exciting technological frontier full of opportunities, and we should be harnessing these opportunities for the benefit of local industries like our creative and media sectors, working with them, rather than at their expense.

This paper works within those parameters. It seeks comment on a scheme that delivers creator protections while creating a viable path for AI training in Australia.

Scope

This paper operates within three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

All three parameters must be met simultaneously. This paper does not consider proposals that achieve one or two at the expense of the others.

This paper takes no position on the character or motives of AI companies or rights holders. It focuses on identifying a workable policy framework.

We invite comments that fall within these parameters. The question is whether a practical arrangement exists that satisfies all three.

Proposal

This paper proposes the creation of a two-part Australian AI Training Scheme – a permit framework for AI training in Australia and a separate fund to support Australian creators.

AI Training Permit

Entities conducting AI training in Australia may apply for an AI Training Permit. Training conducted without a permit must comply with existing laws, including copyright law.

The permit would carry the following obligations:

- **Fee.** Permit holders would pay a fee into consolidated revenue. The fee would be set by Government and could be tiered to reflect differences in purpose and scale. Low or no fee permits could be available for training of limited commercial value, education and academic research.
- **Consent and control.** Permit holders would be required to respect an effective, practical, and free opt-out mechanism for Australian rights holders. The design of the opt-out mechanism would be developed in consultation with rights holders, AI companies, and researchers.
- **Data centre expectations.** Permit holders would be required to comply with the National Data Centre Expectations, including commitments on energy, water, emissions, and community impacts.
- **Other national interest conditions.** The permit could include additional conditions where justified, such as giving notice to the Australian Government of training runs, transparency and disclosure requirements around models created in Australia, and a prohibition on circumventing technological protection measures (TPMs) within the meaning of the *Copyright Act 1968*.

Permit holders would be authorised to conduct AI training in Australia with the legal certainty that such training would not be subject to copyright restrictions (the separate prohibition on circumventing TPMs continues to apply). The privilege attaches only to permit holders in compliance with their obligations. It does not apply to entities that do not hold a permit or that fail to meet its conditions. The Department of Industry, Science and Resources would also provide permit holders with direct assistance in understanding and managing regulatory requirements across local, state and territory and federal jurisdictions.

Australian AI Public Benefit Fund

Government would separately establish and administer a fund, resourced from consolidated revenue, to support Australian creators. The *Public Lending Right Act 1985* provides a tested model.

The fund would serve a dual purpose: compensating Australian creators for income loss, and supporting the enrichment of Australian culture. Eligibility would be limited to Australian citizens and permanent residents who are creators, publishers who publish in Australia and their respective rights holders. The Minister would appoint a committee of relevant stakeholders – supported by staff within the responsible department – to administer the fund and develop operational rules, including how eligible material is identified and how payments are calculated.

The fund's detailed design would be developed in consultation with creators to ensure it meets their expectations. Because the fund is decoupled from negotiations with AI companies, the Government has substantial latitude to design it to meet creator and rights holder expectations.

Policy design

Linking obligations to privileges

The proposal places obligations on those engaged in AI training: paying a fee, giving practical consent and control to rights holders, complying with the National Data Centre Expectations, and meeting other national interest conditions. Only entities that meet these obligations receive the privilege of training AI in Australia under the permit.

A key reason a TDM exception is not a promising policy direction is that it grants a valuable privilege without corresponding obligations in the national interest. Once a TDM is in place, Government would have limited levers to require consent, control, and compensation. A TDM also does not allow for additional nuance on national interest matters beyond copyright.

Any policy will require some linkage between the *Copyright Act 1968* and a broader scheme. Under this proposal, the permit is the vehicle for that linkage. Enabling amendments to the *Copyright Act 1968* would grant privileges to permit holders. Those privileges would apply to all content protected by Australian law.

Asymmetry between Australian and global creators

The fund supports Australian creators. It does not serve global creators. In this regard, it mirrors the *Public Lending Right Act 1985*, which only serves Australian authors and publishers.

The asymmetry is practically necessary. If the fund were to serve all creators or rights holders globally, it would likely be uneconomic.

The permit and the fund are decoupled to manage this. The permit would be available for entities wishing to conduct AI training in Australia, and—via enabling amendments to the *Copyright Act 1968*—authorises training on all content protected by copyright under Australian law. The fund is a separate domestic support mechanism, resourced from consolidated revenue, directed at Australian creators. The two operate independently.

Decoupling also creates policy optionality on both sides of the scheme. The permit fee reflects what AI companies will pay for market access, legal certainty, and accelerated time to power. The fund reflects what Government chooses to invest in Australian creators and Australian culture. These are different questions that might need different scopes, quantum and other policy design variables.

Setting the permit fee

The permit fee would be set by the Government.

AI companies make arguments about the value of training data: that copyright material is a small fraction of overall training data, which also includes non-copyright material, synthetic data, and reinforcement materials owned by the companies; that the value of any discrete data point is negligible; that the value comes from aggregate data; or that Australian data is a small fraction of what AI models train on.

These arguments are not material to the fee. The permit fee is not a purchase price for copyrighted material. It is a fee for the privilege of training AI in Australia under the scheme, which includes the authorisation to train without copyright restrictions and the other conditions of the permit.

Through tiering, the fee design can reflect differences in purpose and scale. Some AI training does not produce a commercially valuable product, such as academic research. Some use copyright material to produce highly commercially valuable products. The fee structure can accommodate these differences.

Design and administration of the fund

The *Public Lending Right Act 1985* provides a model that the Australian AI Public Benefit Fund can follow.

Dual purpose. The *Public Lending Right Act 1985* has a dual purpose. First, it makes payments to Australian creators and publishers in recognition of income lost from their books being available in libraries. Second, it supports the enrichment of Australian culture by encouraging Australians to create and publish books in Australia. The AI Public Benefit Fund could adopt a similar dual purpose. This would meet creators' direct ask for compensation and provide a structural basis to support Australian creative industries amid the disruption AI is likely to cause.

Eligibility. As with the *Public Lending Right Act 1985*, eligibility would be limited to Australian citizens and permanent residents who are creators, and to publishers who publish in Australia. The fund could cover creators across the forms of material that AI models are trained on, including text, images, music, code, and video.

Triggering and calculating payments. The *Public Lending Right Act 1985* pays creators based on estimated copies of their books held in Australian public libraries, calculated from an annual survey. The direct equivalent – identifying which works are held in AI training datasets – is not straightforward. Identifying whether a specific work appears in a training dataset can be technically challenging. The fund would need a workable proxy. Options include using existing institutional infrastructure – such as membership of recognised collecting societies like the Copyright Agency or APRA AMCOS, or listing in established catalogues like the ISBN system. Alternatively or additionally, the fund could support the creative industry more broadly, in recognition of the benefits of the creative industry to Australia as a whole. The appropriate mechanism is a matter for further consultation.

Distribution. As with the *Public Lending Right Act 1985*, payments would be calculated from a funded pool, with a tapering formula to ensure the fund supports a broad base of creators rather than concentrating at the top.

Administration. As with the *Public Lending Right Act 1985*, the Minister would appoint a committee comprising relevant stakeholders – creators, publishers, and other appropriate representatives – to administer the fund. The committee would be supported by appropriate staff within the responsible department. Operational rules, including the mechanism for identifying eligible material and calculating payments, would be set by the committee and approved by the Minister via legislative instrument. This arrangement keeps the fund adaptive as AI's impact evolves, without requiring legislative amendment each time operational rules need to change.

Size of the Government contribution. As with the *Public Lending Right Act 1985*, the Government would decide how much to pay into the fund. The amount is a matter for budget processes and ongoing policy judgment.

Designing the consent and control mechanism

Consent and control, alongside compensation, are responsive to the concerns of Australian creators and rights holders. Any solution needs to meet their expectations while being practical to administer.

Permit holders must offer an effective opt-out for Australian rights holders.

The scheme adopts opt-out rather than opt-in. Opt-in would require AI companies to secure permission from every rights holder whose material might be in a training corpus. This is likely unrealistic.

Three design conditions are essential for the mechanism.

Effective. The approach should account for technical mechanisms that lower regulatory burden (such as machine-readable signals like robots.txt), while ensuring Australian rights holders can effectively opt out across all types of works.

Practical for rights holders and AI companies. An opt-out that requires legal knowledge, fees, or complex identification of works will not deliver meaningful control. Rights holders should be able to opt out simply, at the level of individual works, categories of work, or all of their work. Equally, an opt-out that requires manual tracking of an unlimited number of works would not achieve the goal of facilitating economically viable AI training in Australia.

Enforceable. Permit holders must make reasonable efforts to respect opt-outs as a condition of their permit and be able to demonstrate the processes that they use to ensure compliance. Non-compliance should result in fines that escalate to the loss of the permit. Creators should have a practical avenue to raise concerns.

Technological protection measures are a related signal of non-consent. Where a rights holder has applied a TPM to their material, the permit does not authorise circumvention. Permit holders must comply with Part V, Division 2A of the *Copyright Act 1968*, and breach is a ground for enforcement action under the scheme.

The specific design should be developed in consultation with creators, rights holders, AI companies, and other experts.

Enforcement

Enforcement would be administered by the Department of Industry, Science and Resources (the department), which already administers the National Data Centre Expectations, in coordination with the Attorney-General's Department. The department would have the power to monitor compliance, investigate breaches, and take action against permit holders, including imposing financial penalties and suspending or withdrawing the permit. Grounds for action by the department would include failure to pay the permit fee, failure to respect opt-outs, failure to comply with the National Data Centre Expectations, failure to meet other national interest conditions, or failure to comply with other Australian laws.

Entities that train AI in Australia without a valid permit would not receive the copyright authorisation the scheme provides. They would remain exposed to infringement action by rights holders under the Copyright Act 1968.

Consultation questions

On the overall proposal

1. Does the proposed Australian AI Training Scheme — a permit and fund — represent a workable path to meeting all three parameters in the Scope section? If not, what parameter is not met and why?
2. Are there alternative arrangements that meet the three parameters? If so, what is that proposal, how does it meet the three parameters, and how does it compare to the proposed Australian AI Training Scheme?

On the AI Training Permit

3. The scheme proposes opt-out as the primary consent and control mechanism. What approach would best deliver an opt-out that is effective, practical for creators and enforceable by Government? If you prefer opt-in, how could that be achieved within the parameters?
4. Are there national-interest conditions that should be attached to the permit from day 1?

On the Australian AI Public Benefit Fund

5. Who should the fund support, and how should eligibility be defined? Should the fund focus on creators whose materials have been used to train AI, creators who have lost earnings because of AI, the ongoing viability of Australia's creative industry, or other matters? Why?
6. What should the role of established organisations like collecting societies be in designing, distributing and managing remuneration through the fund to creators?
7. How should compensation under the fund be calculated? What approach best balances feasibility, fairness, and administrative cost?

Open question

8. Are there issues raised by this proposal that are not addressed in the paper but are relevant to meeting the three parameters set out in the Scope section?

This paper invites comment from creators, collecting societies and rights management organisations, AI companies, and the public. Responses are sought within the three parameters:

1. Delivering Australian creator expectations for consent, control, and compensation.
2. Consistency with the Government's commitment not to create a TDM exception, while preparing for future copyright challenges emerging from AI.
3. Facilitating economically viable AI training in Australia, in the broader national interest.

All three parameters must be met simultaneously. This paper does not consider proposals that achieve one or two at the expense of the others.

This paper takes no position on the character or motives of AI companies or rights holders. It focuses on identifying a workable policy framework.

Frequently asked questions**Why is the fund funded from consolidated revenue rather than from permit fees?**

The fund and the permit are decoupled. Permit fees flow into consolidated revenue. The fund draws from consolidated revenue. The scheme would be designed so that there is no direct link between the two.

This design has several purposes. It ensures the fund is not characterised as a de facto licensing scheme, where AI companies pay rights holders for the use of their material. It preserves the Government's ability to set the permit fee and the fund independently, based on different considerations. It also mirrors the Public Lending Right Act 1985, where Government funding is not directly tied to any specific revenue source.

Decoupling also avoids constitutional constraints that arise where a Commonwealth charge is tied to payments to private parties. The fee is a charge for the privilege of training. The fund is Commonwealth expenditure.

Why can't AI companies just license the Australian material they need?

Australian copyright law protects all copyright, not just Australian copyright. The corpus of material protected under Australian law is very large, and most of it is available at low or no cost in jurisdictions with fair use or TDM exemptions. Licensing only Australian material would not authorise AI training in Australia under current copyright law. Licensing all material protected under Australian law would be practically unworkable and could undermine AI companies' legal position in other jurisdictions, particularly the United States.

Why not just require AI companies to pay creators directly through a collective licensing scheme?

Collective licensing is one option being considered by the Copyright and AI Reference Group. AI companies would pay licence fees to a collecting society, which would distribute payments.

The proposed scheme differs in three ways.

First, a collective licensing scheme administered under the *Copyright Act 1968* would cover all material protected under Australian law, not just Australian material. AI companies would pay for a licence covering the global corpus, even though the same material is available at low or no cost in jurisdictions with fair use or TDM exemptions. This would reduce the compensation received by Australian creators while leaving the same practical problem as direct licensing – routing it through a collecting society does not reduce the scope of what is being licensed.

Second, a collective licensing scheme treats the payment as a fee for the use of copyrighted material. AI companies have indicated they are unlikely to accept this, as it could undermine court rulings in other jurisdictions where AI training is treated as fair use.

Third, collective licensing addresses copyright only. It does not provide a vehicle for other national interest conditions – such as compliance with the National Data Centre Expectations, notice of training runs, or transparency requirements.

The scheme does not preclude collective licensing arrangements for other uses of copyrighted material, such as outputs of AI models or non-training uses, or for non-permit holders.

The Government ruled out a TDM exception. Isn't this just a TDM?

No. A TDM exception grants a privilege without a corresponding obligation. This proposal grants a privilege only where specific obligations are met – including payment of a permit fee, practical consent and control mechanisms for rights holders, compliance with the National Data Centre Expectations, and other national interest conditions. The privilege attaches to permit holders in compliance with their obligations, not to AI training generally.

Does the scheme infringe Section 301 of the US Trade Act of 1974?

No. The scheme imposes reasonable and justifiable obligations on a specific activity – AI training – occurring within Australia. It is country-agnostic. It does not discriminate between AI companies based on nationality, and it applies equally to domestic and foreign entities. It does not restrict market access. AI companies that do not hold a training permit can continue to offer their AI products in the Australian market or train AI without the permit's privileges and obligations.

What happens to AI models trained overseas but deployed in Australia?

The scheme regulates AI training in Australia only. It does not restrict the deployment of AI models in the Australian market. Models trained in other jurisdictions – including in jurisdictions with fair use or TDM exemptions – can continue to be offered to Australian users.

The scheme is focused on attracting AI training activity to Australia, supporting Australian creators, and securing the benefits of AI. It does not cover AI use, deployment, or adoption.

Does the scheme affect creators' existing rights under copyright?

No. Creators retain their existing rights under the *Copyright Act 1968*. The AI Training Permit only authorises AI training conducted in Australia by permit holders in compliance with the scheme. It does not affect rights with respect to AI training occurring overseas, training not under a permit, AI outputs, or any other use of copyrighted material.

What counts as "training in Australia"?

Training in Australia is based on the location of the AI training compute. It covers AI training occurring wholly or partly using chips physically within Australia's jurisdiction.

Is the permit fee a tax on AI training?

No. The permit fee is a fee for a specific privilege – the authorisation to train AI in Australia under the scheme, including the copyright authorisation and the other conditions of the permit. It is not a general revenue measure. Entities that do not hold a permit do not pay the fee and can still train AI in Australia, just without the scheme's privileges and obligations.

Does all AI training in Australia need a permit?

No. AI training does not require a permit. Permit holders gain legal certainty that using the data for AI training is not subject to Australian copyright restrictions. It would also provide the permit holder with direct engagement and assistance from the department in understanding and managing regulatory requirements at the federal, state and territory, and local levels. AI training can occur without these privileges and obligations.

Who is Good Ancestors and why is it running this consultation?

Good Ancestors is an Australian not-for-profit. We develop evidence-based policy on issues with long-term consequences for Australians, including AI.

We think AI carries generational risks and generational opportunities. Copyright has become the sticking point in Australia's AI policy conversation. We don't think a workable path has to come at the expense of rights holders or Australia's long-term national interest. This paper proposes one way to do both. We're running this consultation to test it.

Attachment B: What We've Heard So Far

Consultation Paper on an Australian AI Training Scheme

The consultation remains open. Below is an initial summary prepared for the Senate Inquiry into AI and Data Centres. Comments are not attributed to protect the anonymity of participants.

A clear preference from rights holders for licensing

All rights holder groups expressed an unambiguous preference for licensing as the path forward. No rights holder group was open to exploring alternatives.

Lack of detail about licensing as a pathway

We heard claims from rights holders that collective licensing could cover most AI training requirements. This included claims like “80% of all music could be licensed under four deals”.

So far, we have found these types of claims hard to interrogate. No group was able to detail the extent of existing collective licences, the extent of possible extended licences, the extent of reciprocal agreements, and the overall extent of overlap with AI training datasets.

We also found the specific claim of “80% of all music could be licensed under four deals” hard to evaluate. We could not evaluate whether “music” refers to “commercially produced music by musicians working with record labels” or “all music, including user-produced music, music on platforms such as YouTube, and music otherwise not linked to record labels”, or something else. We also don't know whether the “80%” refers to market revenue share or the count of works.

The lack of clarity accounts for part of the disagreement between AI companies and rights holders, and underscores that rights holders and AI companies do not see eye-to-eye on the issue or on how licensing could work.

- AI companies focus on the **quantity of publicly available work** to ensure certainty about their entire AI training datasets. The normal commercial value of a work is largely independent of its contribution to AI training. Music is not a highly valuable data type for AI training.
- Rights holder groups focus on the **commercially valuable works created by the artists** they represent and are less likely to represent low-commercial works.

The large number of copyright-protected works with low commercial value are a focus for AI companies but of less interest to rights holder groups.

Lack of clarity about the starting point for licensing negotiations

We heard claims from rights holders that they have not been approached by AI companies to explore licensing as a pathway.

We heard claims from some AI and technology companies that they have approached rights holders.

We heard claims from other AI and technology companies that they don't see value in negotiations for material that can be readily licensed while there's no solution for the remaining works with no plausible pathway.

Lack of incentive to agree on licensing pathways

Rights holder groups expressed lack of trust, or even animosity, towards AI companies. Rights holder groups questioned the consultation's scope, which asked respondents not to attack character or motives. Rights holder groups argued that the historic behaviour of AI companies is relevant to their character, and said they were reluctant to deal with companies that have a track record of breaking the law.

Rights holder groups also argued that they were on track to win in the US courts, after which they would receive compensation and potentially damages for most AI training. That reduces their incentive to reach an agreement in Australia, which would at best cover some models by some companies.

AI and technology companies said they did not consider licensing to be a viable pathway to training frontier models in Australia.

Given that a deal requires agreement between one or more AI companies and many rights holders and a deal has not been reached to date, we assess that a deal is highly unlikely unless a change occurs.

Disagreement about prospects in US courts

We heard from rights holders that they were confident of success in the US courts. Specifically, they believed claims that AI training has been upheld as fair use are overstated, and that US courts are on track to find that AI training is not fair use. Further, rights holder groups believe that—once they succeed in US courts—the Trump administration would not intervene to favour AI companies. This could lead to them receiving compensation and damages from all AI companies training in the US.

We heard from AI and technology companies that they are confident that US court decisions so far uphold that the act of AI training itself can be fair, but concede that in some cases the circumstances of the training (such as certain modes of obtaining data) were not. AI and technology companies think US courts will continue to hold the view that the act of training itself is fair use.

Further, AI and technology companies believe that the Trump administration would likely intervene to favour AI companies over rights holders if courts found that AI training was not fair use.

Other consultation responses tended to support the view that the Trump Administration was likely to side with AI companies on copyright issues, citing specific remarks made by the President.

The disagreement about how US courts are likely to act is important to understanding the path of possible negotiations. There is further discussion of the most likely alternatives to negotiated outcomes below.

Disagreement about plausible alternative countries

AI and technology companies tended to claim that Australia was a highly desirable location for AI training, but that there were viable alternatives with workable copyright regimes and other desirable factors, such as adequate energy availability. Japan was given as an example.

Rights holders tended to claim that alternatives to Australia were overstated.

Low tolerance for uncertainty from tech companies

We heard a low tolerance for legal uncertainty from tech and AI companies. Companies worried that if they followed a licensing path and secured relevant rights for works controlled by rights holders, there would be substantial residual litigation risk from creators and rights holders not covered by those deals. Their view was that extensive investment required a high degree of legal certainty.

Concerns about the certainty of the fund

Rights holder groups expressed concern about the lack of certainty surrounding the Australian AI Public Benefit Fund. Specifically, they worried that a current or future government may not contribute meaningfully to the fund even if the Permit raised significant funds, or that the fund may fail to distribute money in a way that meets the needs of creators and rights holders.

Rights holder groups also expressed concern that if the scheme fails to attract AI training to Australia, AI companies would never buy a permit, and the fund may never operate, meaning Australia's copyright laws would be weakened for no benefit.

Concerns about Australia's treaty obligations

Some rights holder groups expressed concern that extending opt-out only to Australians may violate treaty or other international obligations.

The consultation paper currently addresses this risk in two ways:

- First, opt-out is a feature of the permit, not of the fund. Australian copyright law would not have differential treatment of Australian works. Instead, any differential treatment would stem from the technical practices of AI companies that hold a permit.
- Second, the consultation paper has optionality to extend the opt-out beyond Australian creators and rights holders.

The Japanese Model

The consultation paper specifically asked respondents who didn't support the AI training scheme to suggest alternative pathways that were not a TDM, would support creators, and would facilitate AI training in Australia.

Some parties mentioned the Japanese model. The original consultation paper characterises the Japanese model as a TDM. We found there was confusion about how the Japanese model operated. The relevant law is Article 30-4 of the Japanese *Copyright Act*. This provision operates as follows:

- A work may be exploited in any way, and to the extent necessary, provided that it is not the "purpose" of the exploiter to enjoy the thoughts or sentiments expressed in the original work (the 'non-enjoyment' requirement); and
- The exploitation of that work does not "unreasonably prejudice" the interests of the copyright owner in light of the nature or purpose of the work or the circumstances of its exploitation (the 'lack of prejudice' requirement)

There is some confusion about the operation of the non-enjoyment requirement. The requirement attaches to the enjoyment of the thoughts or sentiments expressed in **the original copied work**. So, the fact that someone downstream 'enjoys' the model's outputs does not mean that training that AI model was a prohibited use on the basis of 'enjoyment'.

In 2024, the *Japanese Copyright Subcommittee* considered each requirement with respect to the training of generative AI systems.

- With respect to the non-enjoyment requirement, the Committee found that data analysis, which is expressly included as an example of a use fulfilling the non-enjoyment requirement, would generally include the training of generative AI.

- With respect to the lack-of-prejudice requirement, the Committee found that it may not be satisfied in certain circumstances, such as the use of database works created specifically for data analysis, or the use of pirated works.

Article 30-4 contains neither a non-commerciality requirement nor an opt-out for rights-holders. While the Committee notes that each case turns on its own facts, the provision appears to operate largely as a TDM exception, and offers rights-holders less protection than the Australian AI Training Scheme proposal.

[Meta](#), [Google](#), and [Microsoft](#) have referred to Article 30-4 as a TDM in evidence to the UK House of Lords.

What this means for the proposal

The consultation surfaced specific changes that would improve the model. For instance:

- Legislation to create the fund should mandate a committee like the [PLR Committee](#).
 - The new Committee should include representatives from creator and rights holder groups to increase confidence that the fund would be well administered from their perspective.
- Legislation to create the fund could mandate that distributions be made in part through rights holder groups, declared collecting societies or organisations that administer statutory licences.
 - Funds should also be available to creators not represented by rights holder groups.
- Legislation to create the fund could include mechanisms to increase confidence that future governments would contribute. For instance, there could be a specific endowment and an investment mandate, like the [Disaster Ready Fund](#).
 - The DRF is a \$1 billion fund that distributes \$200 million per year.
- The Permit and the Fund could sunset after a given period (e.g., five years) in the event that no permits are ever issued. This would address rights holder concerns about the legacy of the scheme if it is never used.
 - If permits are issued, the legislation should not sunset to ensure certainty for companies that invest in Australia on the basis of the permit scheme.
- Legal advice is required on whether the opt-out mechanism should be global or can be Australian only.

The consultation surfaced opportunities for clarification. For instance:

- Some commentary suggested that the fund would be a once-off payment.
 - The permit is intended to be annual and the fund ongoing.
- Some commentary suggested that the AI training scheme would undercut licensing / data access deals that have already been made.
 - The scheme's intent is to preserve valuable data access opportunities. A permit would not authorise the circumvention of TPMs, such as paywalls. This means commercial deals between AI companies and companies like HarperCollins and the Associated Press involving content behind TPMs would remain on the table.

Overall, the proposal assumes agreement that the status quo is bad for all parties. Specifically, Australian works are currently being used to train AI overseas, with Australian creators and rights holders receiving little or no compensation. Finding a way to attract AI training *and* support Australian creators would be better overall. The proposal also assumed that parties would seek a deal better than the status quo. Under these assumptions, there is a large zone of possible agreement between stakeholders, and the proposal is better than the status quo for all stakeholders.

The main point the consultation surfaced is that **there is little common ground between stakeholders and little appetite to seek mutually beneficial deals**. One explanation for this is different views about the most likely alternative to a negotiated agreement (MLATNA).

Specifically, the paper assumed that the zone of possible agreement was large because parties agreed that the **most likely alternative was the status quo**—that AI training continues overseas where Australian creators and rights holders get little compensation. However, the consultation suggested that different stakeholders have different views on the most likely alternative:

- Rights holders believe that the most likely alternative is victory in the US courts, leading to payments from all AI companies and potentially damages, as well as future licensing deals.
- AI and technology companies believe that the most likely alternative is that they win in the US courts, or that the Trump administration intervenes, or that they train in another jurisdiction with fair use or a TDM exemption.

The below table sets out five possible future scenarios. One is a directly negotiated solution (like licensing). One is a government-negotiated solution (like the AI training scheme). The remaining are alternatives to negotiated solutions based on different outcomes in the US courts and different actions by AI companies. The table illustrates that any negotiated solution is better for Australia's national interest, and that alternative solutions are only better for creators if they prevail in US courts.

Possible future scenarios

The below table sets out five possible future scenarios. A negotiated solution (Scenarios 1–2) is the only path that somewhat secures Australia's national interest. For creators, the alternatives beat a negotiated outcome in just one case (Scenario 3), which requires a US court win and no intervention by the White House.

Negotiated solutions

Scenario	Creators/rights holders	Australia's national interests
1 Directly negotiated solution: AI companies and rights holders agree licences and reciprocal arrangements that allow AI training in Australia	 Australian creators get opt-in and compensation when AI is trained in Australia.	 AI training happens in Australia, but further mechanisms are required to ensure that training happens in the national interest
2 Government-negotiated solution: Australia implements the Australian AI Training Scheme	 Australian creators get opt-out, and compensation via the fund. But exchange a property right for a government fund.	 AI training happens in Australia with a specific mechanism to secure the broader national interest

Alternatives to a negotiated solution

Scenario	Creators/rights holders	Australia's national interests
3 Alternative to a negotiated solution: AI companies double down on US investment. Creators win key court cases in the US, and the Trump administration does not intervene.	 Australian creators get opt-in and compensation when AI is trained in the US.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.
4 Alternative to a negotiated solution: AI companies double down on US investment. Creators lose court cases in the US or the Trump administration intervenes to favour AI companies.	 Australian creators have no opt-out, no credit, and no compensation when AI is trained in the US.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.
5 Alternative to a negotiated solution: AI companies train in other jurisdictions where copyright laws are favourable to them.	 Australian creators have no opt-out, no credit, and no compensation when AI is trained in other jurisdictions.	 AI training does not happen in Australia, meaning reduced economic benefit and reduced ability to secure the national interest.

 The most desirable outcome  An outcome with pros and cons  The least desirable outcome

Likelihood of possible scenarios

How stakeholders approach the negotiated options depends heavily on how likely they think each alternative is.

A directly negotiated solution is very unlikely

Scenario 1—direct negotiations between AI companies and rights holders—is **very unlikely**. Rights holders and AI companies do not agree on key issues, and many have said they don't trust the other parties and/or don't want to negotiate a direct agreement.

A Government-negotiated solution is possible

Scenario 2—Government intervention to create a path or to change incentives or frameworks—is **possible**. Government has the authority to find a negotiated solution. Only a Government-negotiated solution allows Government to protect Australia's broader national interests.

The alternatives to a negotiated solution depend on the United States

The other scenarios (3–5) are what happens if no deal is reached. Which one might eventuate depends on the US courts and US administration. For rights holders, winning in the US courts (Scenario 3) is the best alternative to a deal because it could secure opt-in, compensation, and potentially damages. For rights holders, winning in the US might be better than any negotiated solution.

If rights holders don't win in US courts, Scenarios 4 and 5 are the default.

In every alternative, training happens offshore and Australia's national interests are not secured.

What this means for policy makers

All stakeholders should work to avoid scenarios 4 and 5.

In almost all cases, negotiated solutions are better for all stakeholders than not reaching a deal.

Policy makers, regardless of their priorities, who believe that the US courts and the White House will favour AI companies should prefer a negotiated solution (1 or 2) in Australia.

Not reaching a negotiated solution is only good for policy makers who both prioritise the interests of rights holders **and** believe that rights holders are likely to win in the US courts and receive support from President Trump (Scenario 3).

Policy-makers who agree that a directly-negotiated solution (Scenario 1) is not on track should pursue a Government-negotiated solution (Scenario 2).