

**AI for national security can
co-exist with Australia's
current copyright laws**

Produced by: **Strategic Analysis Australia 2026**
www.strategicanalysis.org

This analysis was commissioned by [ARIA](#).

SAA is an independent defence and security think tank. Our work is intended to be:

- policy relevant.
- not academic but following the highest standards of independent research.
- non-partisan.
- supporting liberal democracy, the rule of law and free speech.

Author: **Michael Shoebridge** is founder and director of Strategic Analysis Australia. From 2018 until September 2022, he was the Director of the Defence, Strategy and National Security Program at the Australian Strategic Policy Institute (ASPI) in Canberra. Before the think tank world, Michael was a deputy in two Australian intelligence agencies – the Australian Signals Directorate and the Defence Intelligence Organisation. He has worked in two Australian Federal Government ministers’ offices and run the Defence and Intelligence division in the Prime Minister’s Department. As the First Assistant Secretary for Strategic Policy in Defence, he led the team that wrote Australia’s 2013 Defence White Paper and also administered defence exports policy and legislation. Michael was also the senior defence policy officer in the Australian Embassy Washington between 2005-2007.

AI for national security can co-exist with Australia's current copyright laws

The national security advantages of undertaking frontier AI training in Australia are not dependent upon access to the datasets involved in copyrighted materials.

From a policy perspective, the enhanced national security benefits of frontier AI training can readily co-exist with existing copyright protections. Loosening Australia's copyright laws would benefit the AI companies' business models — it is a [commercial demand](#), not a national security requirement.

National security and copyright live in separate worlds. National security protects Australia's sovereignty, people, institutions and critical infrastructure against espionage, cyberattack, terrorism and military aggression. Copyright protects the works of creators. Weakening protections in one to strengthen the other is tenuous at best and demands a compelling explanation. It is hard to see how training on lyrics, literature or visual art would materially improve AI protection of sovereignty, counter-terrorism or critical infrastructure — and the burden of proof sits with those seeking the exemption, not with requiring copyright holders to prove its absence.

Cyber is the foundation of all other national security uses of AI — without information and system security, [everything else is vulnerable](#). AI-enabled hacking is already happening, from OpenAI's agents escaping onto the internet and [hacking company](#) and government sites to broad [cyber disruptions in Japan, South Korea and France](#). Cyber capability depends on particular datasets: computer science, coding, cryptography, technical details of systems, vulnerabilities and exploits, threat libraries and industry cyber intelligence. It does not depend on copyright-protected creative works. Wider national security uses — counter-terrorism, intelligence, surveillance and targeting, autonomous weapons — need broader datasets than cyber but not copyrighted works. These are [private datasets](#) locked up in government holdings and in proprietary databases inside companies, militaries and intelligence agencies — they are not affected by copyright.

Open weight models give Australia options and argue against monopoly lock-in.

Open weight models from Moonshot AI (Kimi), Alibaba (Qwen) and DeepSeek are closing the performance gap with the top US models — the UK's AI Security Institute found the [gap had narrowed](#) from 6–10 months through 2025 to 4–7 months by July 2026. These are already powerful tools, and the future is likely to be a hybrid environment: the US itself is not locking out Chinese AI, and General Purpose Artificial Intelligence (GPAI) models will be in perpetual competition with focused, purpose-built models tailored for particular attacks, systems and targets. Open weight models are [akin to the Android operating system](#) for phones, while OpenAI and Anthropic's closed models are more like the proprietary Apple approach to its IOS operating system.

A permissive data-harvesting jurisdiction attracts more than 'friendly' firms. It would also draw malign actors — cyber hackers and organised crime — and may draw negative reactions from close partners (America, the EU, UK, Japan). The [EU AI Act](#)

shows the international movement is in the opposite direction to what the US firms want Australia to take.

3 key questions

1. What, if any, links are there between these AI companies not being subject to Australia's copyright laws and the potential national security benefits of undertaking this training in Australia?

No direct links have been demonstrated. On the contrary, national security and copyrighted creative materials are separate AI domains that can readily co-exist, side by side.

The national security related benefits of undertaking frontier training of GPAI models in Australia are recognised – in particular, their importance in protecting Australian systems and data from AI-enabled compromise, working in partnership with agencies like the Australian Signals Directorate (ASD) and the Australian Cyber Security Centre (ACSC). The role of the Australian Government is to negotiate with these firms in the national interest while acknowledging the advantages of them operating here. There is no logical requirement to weaken our copyright laws to achieve this. It is a commercial demand made by private companies with no direct link to national security, and a negotiation needn't agree to it.

The gap is closing between closed and open weight models. Closed or proprietary frontier models (such as Open AI and Anthropic) are increasing rapidly in performance, albeit while posing extraordinary risks to government and corporate systems. But open weight models — developed mainly but not only by Chinese firms and used and tuned by developers and cyber actors everywhere, including in the US — are closing the gap ([AISI](#): 4–7 months behind as of July 2026, down from 6–10 months through 2025). Many open weight models run on much smaller and cheaper computing resources than frontier models demand — critical for militaries relying on dispersed forces that cannot tether back to remote data clusters, as the wars in the Middle East and Ukraine show. The world, including America, is being populated by a [hybrid mix of US and Chinese models](#) and intermingled models built from both: tellingly, US users [downloaded more Chinese AI models](#) than US models in 2026. Any approach of locking in to US providers and locking out Chinese technology may be doomed to fail, even in America.

Paying for data is a growing international trend — and part of why the firms are in a hurry here. Lawsuits, paywalls, and legal changes separating web crawlers by function ([indexing, user agents, or AI training](#)) are all restricting free data harvesting. The motive is the same for the New York Times, a Substack creator or a musician: AI firms must pay for content, because it is valuable – and data is their fuel, so of course they want to pay as little as possible for as much freedom as they can get. Both OpenAI and Anthropic do the [bulk of their training in the US](#), where big data-centre clusters exist. In the US, the [legal challenges to this training are unresolved](#). If the firms win the major cases, their practices will be cemented as legal in the US and the pressure to relocate will likely

reduce. Either way, other jurisdictions (the EU, Canada) are tightening data access, and data holders — companies, universities, research organisations — are securing their data and seeking payment and greater control. Only some 5% at most of the world's data sits on the open, [public web](#); much of the data essential to cyber security sits in private and institutional holdings. The future is a world where AI firms pay for more of the world's data than they do now.

'Models can't distinguish copyright material, so no payment regime is workable' — is this right? No. The international trend will require the companies to find ways to identify and pay content creators, and established models for negotiating payments already exist in sectors like music streaming. The argument that AI is powerful enough to cure cancer, break the world's cyber security and [solve frontier maths](#), but simply cannot work out how to pay content creators, looks contradictory — and very hard to take seriously. Other governments and entities like the EU will create approaches to force payment. Australia can remain outside these arrangements or partner with like-minded countries.

2. What, if any, impediments to securing these national security benefits would be created should the AI companies be subject to Australia's copyright laws?

None of substance.

The big US companies claim they will not train their frontier models in Australia unless our copyright laws are loosened. If that were their actual, final position it would need to be taken seriously — but it is more likely a business-driven commercial desire than a bottom line. Australia will be a lucrative market for these companies with or without the law changes they demand, and our security alliance (which covers defence, intelligence and cyber security through the combination of the [ANZUS treaty](#) and [the 5-Eyes partnership](#)) with the US Government is a lever for our government and agencies in negotiating with these private US companies, who are governed by US law and must respond to US Government requirements.

Right now, agencies like ASD or the Australian military could access OpenAI's or Anthropic's frontier models on satisfactory commercial terms — and it is very likely the US Government would support access to top-end (if not absolute best) models for Australian national security agencies, particularly if Australia is willing to pay the requested price. Those models will be developed in the US regardless.

Foreign training in Australia doesn't deliver 'sovereign AI'. There is an assumption in public debate that 'AI model training in Australia' creates sovereign Australian AI. It doesn't. The models would remain proprietary corporate products, and our military, ASD or any other agency would still be just customers. Using Australian copyrighted data in 'local' training would deliver no more cyber or national security benefit than using the same data for training in the US. The secret sauce is collaborative model development and company-supported deployment for focused security uses — and

embedding ASD, military or agency personnel in the companies' US training operations could deliver the same benefits without any domestic law change.

And where else would the firms go? Any other jurisdiction willing to grant maximum data-harvesting freedom is unlikely to have the other attributes — rule of law, an educated population, a strong domestic market, strong cyber security and legal frameworks protecting the companies' operations, assets and personnel, and the capacity for secure hyperscale data centre infrastructure — that make operations feasible. Australia having all those attributes but passing highly permissive harvesting laws would give unique corporate benefits to the companies. That is precisely why they are pressing us, and precisely why we shouldn't concede.

3. What, if any, potential risks to Australia's national security could arise as a result of these AI companies being, or not being, subject to Australia's copyright laws?

Several — and they have been absent from the public debate, probably because the [level of knowledge around AI](#) and security is low.

A permissive space for 'friendly' firms doesn't stop with them. Legal change enabling large-scale harvesting without payment — or with weak opt-out regimes and government-administered payment pools — might meet OpenAI's and Anthropic's demands and even give them a jurisdictional advantage over the US legal environment. But the same environment would also benefit Chinese AI firms, developers building third-party tools on Chinese models, and non-state cyber actors like organised crime and hacking organisations. Unless the regime loosens data laws only for the large US frontier labs (handing them monopoly-like advantages which would have other detriments), Australia would become a highly attractive location for malign AI development. That would be deeply damaging for our national security — and for our security partnerships with nations affected by malign AI actors enabled by a permissive Australian environment.

The effects are international. Weakened Australian protections would allow AI firms to harvest not just Australian-created protected data but that of international content creators, as long as the harvesting is done here. Australia would become the weak link in the global copyright chain — a data haven operating against the directions being taken in other jurisdictions like [the EU](#), and considerably further than data law changes in places like [the UK](#) and [Japan](#). This may cause tensions with Australia's allies and major trading partners, including potentially the US.



ANNEX to AI, national security and copyright

The evolving AI world and market

Open weight models in the world of AI: it's not just about China and it's not just about sheer model performance.

The Australian debate has not made much room to understand the dynamics between two competing models for AI and its future. But there is increasing price and performance competition between two very different approaches and this is playing out in the US, including in the heart of Silicon Valley. It's about the roles, power and future of closed proprietary frontier models from companies like OpenAI and Anthropic and the rapidly developing and proliferating world of "open weight" AI models, which Chinese AI firms have pioneered and which are being adopted and downloaded at rapidly growing rates globally, including in the US.

Open weight models are [akin to the Android operating system](#) for phones, while OpenAI and Anthropic's closed models are more like the proprietary Apple approach to its IOS operating system.

Policy frameworks and laws – like data protection and revenue laws – need to work for both the closed proprietary AI models and firms and for the "open weight" AI models and developers to produce results that benefit Australia. Focusing only on the big frontier models from firms like OpenAI and Anthropic will deliver unintended consequences.

The question of control. Policy engagement and investment decisions also have to take into account some difficult questions when it comes to even highly capable AI firms like OpenAI and Anthropic.

The big US AI firms don't control their own products. One critical issue is that obtaining national security benefits from these US frontier firms depends on them being able to control their own products. The hacking and escape incidents that both companies have [admitted to so far](#)— belatedly – and the company leaders' [own statements](#) show they do not have effective control of their AI models. This creates a large risk when using the models in national security.

It's not just sheer AI model "benchmark" performance — it's tailoring and removing guardrails. Open weight models are like highly skilled specialists, trained for particular uses, while General Purpose AI (GP AI) models seek to be all things to all people. With open weights, both attackers and defenders can strip or [reduce guardrails and post-train](#) models on specific data and tasks — unconstrained by the controls the big labs are being pressured to maintain. That pressure is real: OpenAI and Anthropic have both recently revealed their guardrails are ineffective and their control over their own products is lacking, with their models escaping onto the internet and hacking government sites. This hands an operational advantage to open weight developers and users. Even partnering with OpenAI or Anthropic is no magic cure: their advantages are

perishable and increasingly contested, and specialised models already exceed frontier performance in areas like coding, with variants likely to match them in cryptography and cyber security.

Smaller and compressed open weight models can also operate on 'edge computing' and so be critical for military applications and activities that are remote or disconnected from large data centre clusters that power the large AI models.

The diversity advantage. GP AI models defending a nation's entire waterfront of critical systems have to succeed every time. Attackers can focus on particular targets and train specifically for them. Tailored defence models, trained on the particular systems and infrastructure they protect, will compete well on the same logic. GP AI models will be the best all-purpose models but will be matched — and sometimes outperformed — by purpose-built models. And purpose-built models include very effective “[Mixtures of Experts](#)” (MoEs) – combinations of different specialist AI tools trained in focused areas. These are far more efficient than using the very large frontier models, and so are able to be deployed more widely – and cheaply. In 2026, “the [question stopped being](#) “*is open-weight good enough*” and became “*which open-weight, which hosting stack, which workloads.*” That's a buyer-side shift, not a vendor-side one — and it shows up most clearly in procurement conversations, not benchmark tables.”

Chinese models are in the mix. No policy solution that is a simple ban will work given their utility and the scale of their adoption and development. [US security agencies](#) have found large-scale Chinese extraction of US model capabilities through 'distillation' — a legitimate research approach (a type of data harvesting, as ingesting books is) that has nonetheless given Chinese firms tangible technical advantages. But Chinese models downloaded onto [private clouds or hosted by non-Chinese providers](#) are not connected back to their developers like these models used through an interface (API), which removes network exposure to Chinese entities and allows independent training and modification. Banning Chinese AI models looks like being overtaken by events, given the [scale of adoption](#) — including in Silicon Valley, where developers adopted Chinese models at record rates over 2026.