

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 leverage 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.

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.