Kate Chaney MP
FEDERAL MEMBER FOR CURTIN

AI DISCUSSION PAPER

Shaping the future of AI for Australia

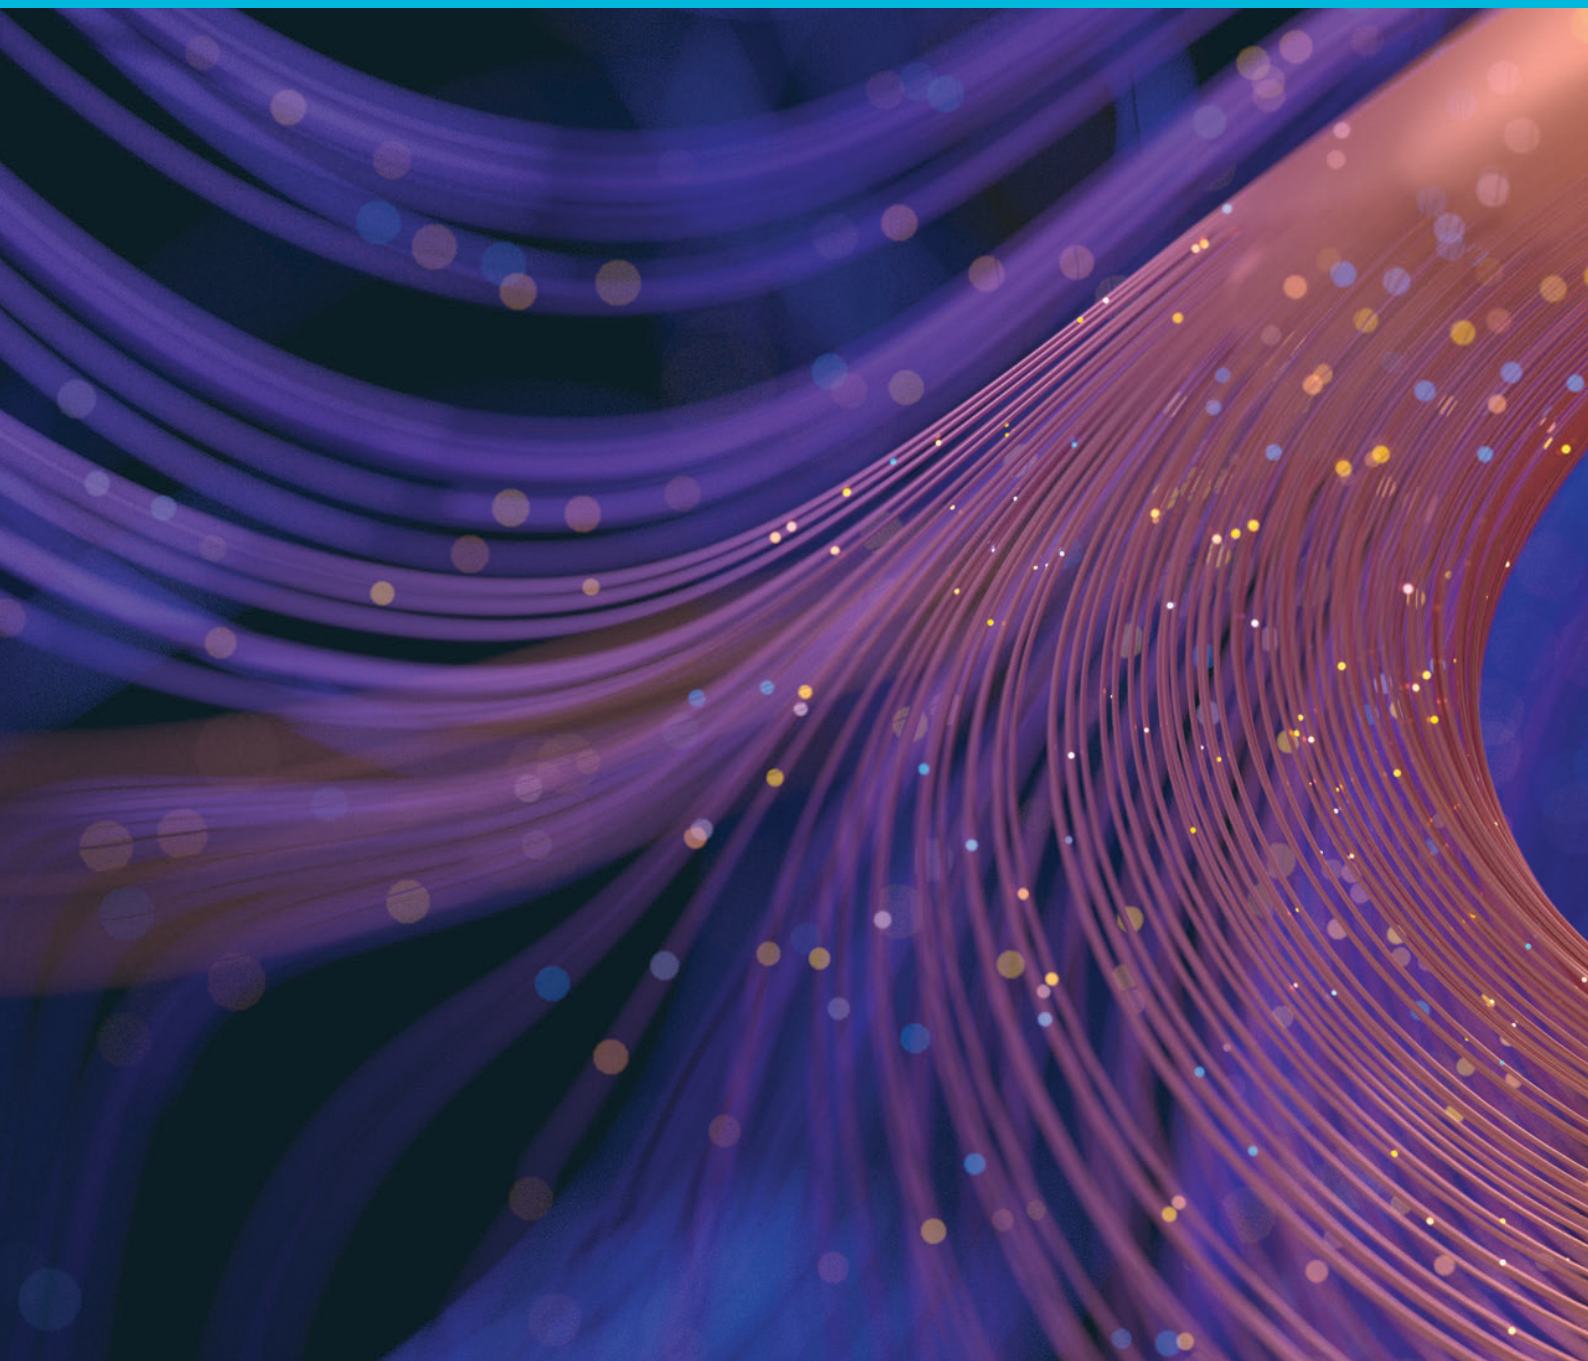

Shaping the future of AI for Australia

Artificial intelligence (AI) is already reshaping economies, labour markets and the information environment. Its ultimate trajectory remains uncertain, but the scale of potential change is comparable to the industrial revolution. The decisions being made right now will shape what that change looks like for Australians.

Australia has done remarkably little to prepare. The Australian Government has identified broad aims for AI policy but has implemented very little actual policy to achieve them. This reflects the genuine uncertainty around AI's future development – no one knows exactly how powerful AI will become, or how quickly. Perhaps the current wave of excitement will ultimately produce little more than a generation of very effective chatbots. But most think the implications could be far larger.

Uncertainty is not a reason for inaction. It is precisely why governments must begin preparing now. A passive, hands-off approach does not mean Australia avoids the consequences of AI – it means **our future is being determined by overseas technology companies** and the billionaires who run them. The Australian Government is effectively hoping that the AI models developed overseas will increase productivity and deliver other benefits without causing major societal disruption. It is doing little to make that outcome more likely and little to cushion Australians if it does not come true.

> "This is one of the most massive technological changes in human history – we need to put a lot of thought into managing it."
>
> **Curtin constituent**

Curtin community survey on AI

Across Curtin, 386 constituents filled out a community survey with their perspectives on AI. Their responses are threaded throughout this paper. Some key results are:

- Curtin residents are 4 times more likely to be concerned than optimistic about AI;
- 8 in 10 Curtin residents do not think the government is doing enough to deal with AI; and
- 8 in 10 Curtin residents do not think that current policy will ensure AI benefits Australians, rather than overseas tech companies.

In this context, it is no wonder that we are seeing a backlash against AI. Australians are worried that this new industry will steal jobs, supercharge online deepfakes and scams, breach their privacy, and use up energy, land and resources – all for the profit margins of international AI companies. Australians do not have confidence that the benefits will be shared with all Australians, because the government has not put in place any policy to ensure this.

The competitive context in which AI is being developed – between companies and countries – makes it feel like a runaway train. But we must also race to govern the technology well, not just to develop it. AI is the only tool ever invented that can make and implement its own decisions. This creates a range of novel risks. Australians expect government to put frameworks in place to keep them safe as we navigate a rapidly changing world.

This paper is an attempt to seize the reins of the AI debate. It asks a simple question: what kind of AI future do we want and what should Australia do now to build it? It presents a set of specific, practical policies that can and should be implemented now to:

- set up the **structures**;
- capture the **opportunities**;
- deal with **current harms;**
- prepare for **emerging risks;** and
- **share the benefits** of AI.

For example, this paper contains actionable approaches the government should be taking now on issues that concern Australians, like **deepfakes**, **unhealthy relationships** between children and AI chatbots, improving **research and productivity**, and **taxing AI companies** so we can all share in the economic benefits.

It is not an exhaustive account of everything governments will eventually need to do as AI develops, but rather a clear set of actions that should not wait. These policies are designed to ensure that Australian voices, not overseas technology companies, are shaping our future.

I would like to thank my constituents in Curtin who completed our community survey, some of whom are quoted in this paper. I would also like to thank the many experts across the country who shared some of their incredible knowledge with me.

Kate Chaney MP
May 2026

"The federal government is almost totally reactive in response. It is completely achievable to be proactive at this time."

Curtin constituent

"The benefits and the dangers are extreme. We have to move fast with policy while we can, before it's too late."

Curtin constituent

"We need to be bold into this next decade. There are risks, but let's lead."

Curtin constituent

Policy priorities

Setting up the structures to govern AI policy

1. Increase funding and resourcing for the **Australian AI Safety Institute**, so it can keep pace with developments in AI and effectively manage risks as they emerge.
2. Increase funding, resourcing, and the mandate for the **National AI Centre**, so it can actively identify and pursue AI opportunities that benefit Australians.
3. **Stress test and update Australian laws** to ensure they are fit-for-purpose in the age of AI.

Capturing the opportunities of AI

4. Establish a **National AI Missions program** to actively identify and pursue a small number of AI opportunities where Australia can lead and deliver real benefits.
5. Ensure that Australia's **copyright system** fairly compensates Australian creators for their content, while giving AI developers the opportunity to invest and operate in Australia.
6. Establish a dedicated **AI for Science program** to ensure Australia's world-class research sector is making the most of the transformative potential of AI.
7. Invest in developing **world-class Australian AI researchers, engineers and safety experts**, and create pathways to attract the best international AI talent to Australia.

Dealing with current harms of AI

8. Establish a **digital duty of care** which places an obligation on digital platforms, including AI chatbots and social media sites, to take reasonable steps to prevent harm to users.
9. **Extend Australia's under-16 social media ban to cover AI chatbots** and companions that simulate emotional intimacy, to protect the psychological and social development of our children.
10. Address the overwhelming volume of **AI deepfakes, misinformation and disinformation**, so Australians can distinguish what's real and what's not.
11. Strengthen **privacy laws** by preventing use of personal data that is not fair or reasonable.
12. Establish a legislated and mandatory **automated decision-making framework** that ensures transparency and safeguards around automated decisions made by government.

Preparing for emerging risks of AI

13. Undertake rigorous **scenario analysis of the potential trajectories of AI's impact on the workforce and economy**, and develop contingency plans so that government is prepared to respond – whatever happens.
14. Establish a **mandatory transparency and incident reporting scheme**, requiring frontier AI developers to publish information on how their models are built and on any serious incidents they cause, so policymakers and the public can assess the risks they pose.
15. Update Australia's **crisis management framework** to include severe AI risks and undertake crisis simulations to ensure government is prepared to respond to catastrophic AI scenarios.

Sharing the benefits of AI with all Australians

16. Establish **national AI literacy programs** so that every Australian, from primary school to the workplace, has the knowledge to use AI well and to protect themselves from its risks.
17. **Convert the government's Data Centre Expectations into binding obligations**, so that data centre development in Australia delivers genuine benefits for Australians.
18. **Tax international technology and AI companies** to ensure that the enormous economic value generated by AI companies in Australia is shared with Australians.

Setting up the structures to govern AI policy

Good AI policy requires institutions capable of keeping pace with AI. Governments move slowly; AI moves fast. We do not know exactly how AI will develop, what new risks will emerge, or what opportunities will arise.

What we do know is that Australia needs structures in place so that we can identify and respond to new developments as they happen, rather than scrambling to catch up after the fact. Regulation can take years, but the latest AI upgrades can be installed in minutes. Without dedicated, well-resourced bodies that understand the technology and can advise policymakers in real time, Australia will always be reacting to problems after they have already caused harm or chasing opportunities after they have already passed. The policies in this section establish the structures that everything else depends on.

"Get ahead of the game – don't try to catch up later, because you simply won't be able to."

Curtin constituent

"They need to create a new paradigm for how they work with and deal with AI. Don't just make a law that will be outdated or redundant in the next 24 hours."

Curtin constituent

Policy 1: Increase funding and resourcing for the Australian AI Safety Institute

Increase funding and resourcing for the Australian AI Safety Institute, so it can keep pace with developments in AI and effectively manage risks as they emerge.

Background

In November 2025, the Australian Government announced the formation of the Australian AI Safety Institute (AISI). The AISI will be a team of AI experts within the Department of Industry, Science and Resources that **helps the government keep pace** with rapid developments in AI, particularly as new risks and harms emerge. It will also work with the international network of AI safety institutes in other countries. The government is still setting up the AISI, having initially planned for it to be operational by early 2026. It will be a vital part of Australia's response to AI. It is very important for policymakers and regulators to have a source of AI expertise that they can draw on when they develop AI policy. It is also important that the government has the ability to keep up with the rapidly changing developments in AI. Governments are not known to move fast, but new generations of AI are being rolled out rapidly – an effective AISI would ensure they can keep up.

Problem

The AISI has not been adequately funded to be able to deliver its important functions. A well-resourced AISI would:

- **test new AI models** to identify safety risks (e.g., it would test whether the latest AI model could help someone build a biological weapon, or encourage a child's plans for suicide);
- **monitor risks from increasing adoption and deployment** of AI tools (e.g., it would monitor the trends of unhealthy relationships between children and chatbots);
- **advise policymakers and regulators** on effective and proportionate AI policy (e.g., if a risk is identified, like a new AI model with unparalleled hacking abilities, the AISI could work with the appropriate regulator or department to develop a policy that will deal effectively with the issue, without being heavy-handed);
- **work with international AI developers** to identify and manage risks (e.g., the AISI would secure access to pre-release models and work with AI developers on solutions for identified risks); and
- **work with international AI safety institutes** to share findings on risks and solutions.

This work is difficult and essential. The AISI is our frontline response to existing and emerging AI harms and risks. Yet Australia's AISI is allocated less than $8 million a year. By comparison, the UK's equivalent organisation, the UK AI Security Institute, receives 15 times more funding. As a result, the UK AI Security Institute is widely respected as the 'gold standard' for such organisations, bringing significant credibility and investment to the UK. Without increased funding and sufficient technical expertise in AI, there is widespread concern that **the Australian AISI will not be able to deliver its vital functions**.

The AISI has other potential issues that must be monitored. There is a risk that it does not have enough technical expertise in AI. There is also the potential that it will be caught up in the bureaucracy and politics of government, without the agility to keep pace with AI and provide long-term, evidence-based advice. Its success also depends on early access to AI models and information, which may require legislation to ensure AI developers are transparent (Policy 14).

Solution

Significantly increase funding for the AISI, so it can effectively manage AI risks as they emerge. For example, a commitment of $100 million each year – a small expenditure compared to many government programs – could:

- ensure the AISI can carry out its complex and vital functions listed above, such as ensuring Australian policymakers and regulators are developing effective and proportionate policy to safeguard Australians from the emerging risks of AI;
- enable the AISI to compete with Silicon Valley and London to attract technical talent; and
- position the Australian AISI alongside the UK AI Security Institute as the 'gold standard' in AI safety, which would boost Australia's credibility and position Australia as a place for the AI industry to do business.

Policy 2: Increase funding, resourcing, and the mandate for the National AI Centre

Increase funding, resourcing, and the mandate for the National AI Centre, so it can actively identify and pursue AI opportunities that benefit Australians.

Background

The National AI Centre (NAIC) was established in 2021 to support Australian businesses to adopt AI responsibly. It is housed within the Department of Industry, Science and Resources and currently provides voluntary standards and practical guidance for businesses. The NAIC is a vital part of Australia's response to AI and an important partner to the AISI. Instead of focusing on identifying and managing risks, as the AISI does, the NAIC focuses on identifying and capturing opportunities.

Problem

The NAIC is too small and too narrowly mandated to fulfil this role. It currently receives less than $9 million a year. This is enough for a small team producing guidance documents and hosting events, but nowhere near enough to identify key opportunities and drive the growth of one of the most important industries of the 21st century. As a result, Australia's approach to AI opportunity is passive. We hope that AI investment flows our way, we hope that productivity improves, and we hope that medical and scientific research flourishes – but we are not doing anything to ensure that this is the case.

Solution

Significantly increase funding and the mandate for the NAIC, so it can actively identify and pursue AI opportunities that benefit Australians. For example, a commitment of $100 million each year – a small expenditure compared to many government programs – could support a range of expanded functions. The expanded NAIC should do the following.

- **Run the National AI Missions program** (Policy 4) – the NAIC should develop a structured process to identify and pursue around five high-value AI opportunities for Australia. These missions would be selected for Australian advantage (where we have strong research, data, expertise or industry capability), national importance (like health, food security, disaster resilience and clean energy) and viability. For example, the NAIC could develop National AI Missions for opportunities such as natural disaster prediction and prevention, wheat crop yield and disease prevention, and clean energy development.
- **Actively identify and work with government on capturing AI opportunities** – beyond National AI Missions, the NAIC should work to identify emerging AI opportunities and collaborate with the appropriate policymakers and regulators to develop the right policy settings for these opportunities to be captured. For example, the NAIC could assess the value of developing data centres in Australia and advise government on policy required to unlock and share this opportunity with Australians.
- **Scale up support for Australian small and medium businesses** – the NAIC could grow its current support for small and medium businesses across metropolitan, regional and rural Australia. The current AI Adopt Centre network should be expanded to cover every state and territory with sector-specific capability.
- **Actively identify and unlock opportunities to share the benefits of AI with Australians** – alongside supporting the growth of the AI industry, the NAIC should focus on opportunities to ensure the economic value and societal benefit of AI are shared fairly with Australians, rather than exclusively captured by international AI companies.

> "Develop domestic capabilities at every level – data centres, local AI models, application into organisations and investment into Australian tech small businesses."
>
> **Curtin constituent**

Policy 3: Review and update key laws to ensure they are fit-for-purpose

Stress test and update Australian laws to ensure they are fit-for-purpose in the age of AI.

Background

Over the last few years, there has been a debate about the best way to deal with the risks presented by AI – do we establish an overarching AI Act or update existing laws, so they cover AI? With the publication of the National AI Plan in 2025, the government decided that updating existing laws was the best approach. Since making this decision, however, there has been almost no publicly visible activity on how the government is approaching this process of reviewing and updating our laws.

Problem

AI presents a range of new risks that are not well addressed by existing legislation. Some areas include the following.

- **Child sexual abuse material (CSAM)** – Australian law prohibits holding CSAM, but does not criminalise the training and use of AI tools designed specifically to generate CSAM. This gap in the law increases the scale and severity of CSAM and makes it harder to locate both victims and perpetrators.
- **AI-generated misinformation and disinformation** – our laws were not built to manage the huge volumes of realistic misinformation and disinformation that can now be produced by AI tools. Particularly in an election context, this gap provides an avenue for foreign interference and large-scale disinformation campaigns, which we are already seeing overseas.
- **Legal liability** – our legal system is generally built on the principle that humans are responsible for their actions and decisions. The question of legal liability is significantly complicated by AI. Now, people may unwittingly and unintentionally break the law through the use of AI agents. Liability for these scenarios should be clarified.
- **Critical infrastructure** – Australia's critical infrastructure, such as infrastructure for energy, water, food, transport, health and financial services, is regulated by the Security of Critical Infrastructure (SOCI) framework. This framework requires operators to prepare and plan for key risks to our critical infrastructure. However, the recent independent review of the SOCI Act found that the regime needs to be modernised for the AI era. Data centres hosting AI models are not effectively covered under current definitions and there is no pathway for covering AI models or services themselves once they become critical infrastructure. Further, the risks posed by AI are not included. The growing capabilities of AI, most recently demonstrated by the Mythos system, are an increasing risk to our critical infrastructure that we are not prepared for.

There are countless other areas of legislation that need updating to be fit-for-purpose in light of emerging AI technologies.

Solution

Stress test and update Australian laws to ensure they are fit-for-purpose in the age of AI. This process should first involve identification of the most likely and consequential AI risk scenarios, followed by an assessment of whether existing legislation is fit-for-purpose. This should be coordinated by the Department of Industry, Science and Resources or the AISI, and should cut across a range of government departments. It should also include a consultation process, so a wide range of stakeholders can identify areas that need clarification or updating.

The review should prioritise areas where there are clearly problems, such as those identified above. There should also be an ongoing process whereby the AISI, existing policymakers, and regulators can identify and deal with gaps in existing laws as they emerge over time.

> "Technology is rapidly accelerating beyond our legislation."
>
> **Curtin constituent**

Capturing the opportunities of AI

The benefits of AI are not guaranteed. Productivity growth, scientific breakthroughs, new industries and better public services are possible, but they will only materialise if the Australian government creates the conditions to encourage private sector innovation.

Hoping that AI investment flows our way, or that overseas technology companies happen to build tools that suit Australian needs, is not a strategy. The policies in this section are about actively capturing the opportunities of AI: identifying where Australia has genuine advantages, funding specific programs to realise them, building the infrastructure and skills base that AI adoption requires, and removing barriers to AI opportunities and unnecessary red tape facing businesses.

In many ways, all the policy recommendations in this paper are about capturing the opportunities of AI. The biggest barrier to AI is trust. Effective and proportionate policy to manage the risks of AI is essential to capturing the opportunities.

"AI is a powerful tool that can unlock significant national benefits. Finding the right balance to not over-regulate whilst encouraging its use is critical"

Curtin constituent

"We will see the end of poverty. New physics. Solve all problems. Eradicate all diseases. The opportunities are boundless."

Curtin constituent

Which AI benefits are you most optimistic about?

Curtin community survey results; respondents picked up to three options

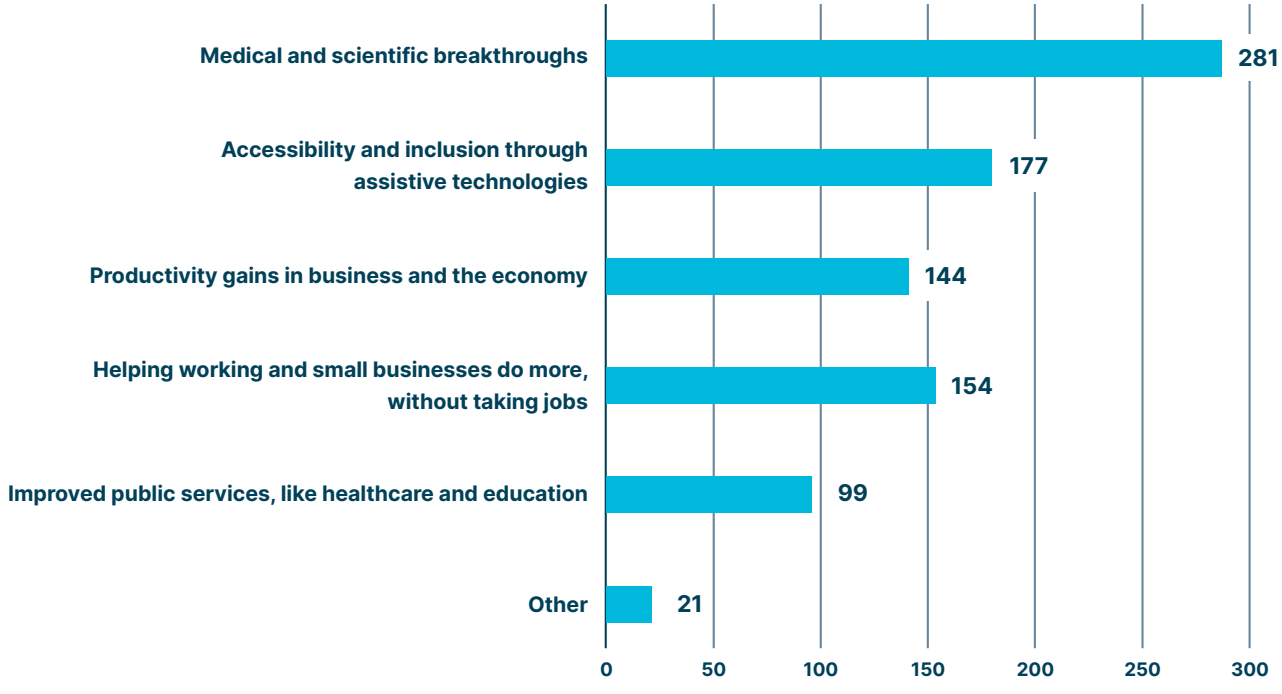

Policy 4: Establish a National AI Missions program

Establish a National AI Missions program to actively identify and pursue a small number of AI opportunities where Australia can lead and deliver real benefits.

Background

The potential benefits of AI are enormous. AI could transform medical research, accelerating the discovery of new treatments and enabling earlier diagnosis of diseases. It could lift productivity across the economy, helping businesses and workers do more with less. It could improve the delivery of government services, making them faster, cheaper, and more accurate. It could help Australia manage the environmental challenges that define our continent – bushfires, drought, reef degradation, climate adaptation. These are huge opportunities – and evidently worth pursuing.

Problem

For now, most of the opportunities of AI remain a promise, rather than a reality. There are genuine, proven examples of AI delivering transformative results. AlphaFold has predicted the structure of over 200 million proteins, fundamentally changing the pace of drug discovery. AI tools are helping Australian farmers detect crop disease earlier and predict yields more accurately. But there are also hundreds of AI announcements that have not yet delivered on their initial promise. The widely predicted productivity boom from AI adoption has not yet materialised in the economic data at scale. The gap between what AI could do and what it is actually doing remains large.

This gap will not close by itself. Countries and sectors that are realising genuine AI benefits are doing so because they made deliberate choices about where to invest and what problems to direct AI toward. Without a structured approach to identify and pursue AI opportunities, Australia risks spreading its effort thinly across too many areas, remaining a follower in all of them. Instead, we should build solutions that reflect Australia's specific strengths and serve Australia's specific needs.

Solution

Establish a structured National AI Missions program to actively identify and pursue a small number of AI opportunities where Australia can lead and deliver real benefits. This would be run by the NAIC. Each mission would be selected against three criteria:

- **Australian advantage** – where we have the research, data, expertise or industry capability to lead;
- **national importance** – where success would deliver real benefits to Australians; and
- **technological viability** – where AI can realistically and significantly contribute to the outcome.

Potential missions include:

- **natural disaster prediction and response** – building on our world-class natural disaster modelling, data and expertise to develop AI tools that give communities and emergency services earlier warning and better information during fire, flood, drought and cyclone events;
- **agricultural productivity and resilience** – applying AI to Australia's agricultural science expertise to improve crop yield forecasting, detect disease earlier, reduce input costs for farmers and help agriculture adapt to a changing climate;
- **health AI built on Australian data** – using Australia's Medicare and PBS datasets, with appropriate privacy safeguards, to develop AI tools that improve disease prediction, diagnosis and treatment for the Australian population; and
- **clean energy transition** – using AI to make the most of Australia's potential as a clean energy superpower by optimising Australia's electricity grid and identifying opportunities for critical mineral and green export projects.

Each mission should be backed with the tools needed to deliver. For researchers and businesses working on mission-aligned problems, this means coordinated access to grants and research funding, R&D tax incentives, support for startups, and priority access to sovereign compute. The goal is not to pick winners arbitrarily, but to concentrate Australia's existing capability and investment behind a small number of areas where we are best equipped to lead and where the benefits of doing so will be felt by Australians.

Policy 5: Use the copyright system to compensate Australian creators and unlock AI training opportunities

Ensure that Australia's copyright system fairly compensates Australian creators for their content, while giving AI developers the opportunity to invest and operate in Australia.

Background

There is an ongoing global debate about whether training AI models on copyrighted material should be considered an infringement of copyright. AI models are trained on vast datasets of text, images, music and other content scraped from the internet. Much of this content is protected by copyright law in Australia and other countries.

Many AI developers argue that their use of the copyrighted material should not be covered by copyright law. They argue that they are not directly copying and reproducing material, rather their AI models are learning patterns from the content – in the same way that humans do. They argue that they should not have to pay compensation to the writers, musicians, journalists, artists, and other creators whose material they use for training.

Conversely, many creators argue that training AI models on copyrighted material should be covered by copyright law – and that licensing agreements should be negotiated between AI companies and rights holders. They argue that the AI companies are building significant economic value from the copyrighted material and that the AI models often directly reproduce copyrighted material when asked.

In Australia, copyright law clearly covers AI training, meaning that AI companies would have to develop licensing deals with copyright holders to use their material for training. In the US, where most of the training occurs, the law is more ambiguous – and is currently being contested in the courts. In some jurisdictions, like Japan and Singapore, AI training is not covered by copyright law.

Some AI companies have argued for a text and data mining (TDM) exception in Australia, which would give them the ability to train on all copyrighted material in Australia without paying compensation. Many Australian creators pushed back against this, and in October 2025, the government ruled out a TDM exception. Instead, it committed to 'fair remuneration' for creators through a licensing model. The government has established the Copyright and AI Reference Group to design that mechanism.

Problem

This is currently a lose-lose-lose situation.

- **Australian creators lose** – because some AI companies are slow or unwilling to arrange licensing deals to use their work in exchange for fair compensation and agreed terms. Worse still, when overseas jurisdictions have more relaxed copyright laws, it means that AI companies can train on Australian copyrighted material for free – as long as they are based overseas. Australian copyright law only protects copyrighted material from AI training in Australia.
- **AI companies lose** – because they cannot undertake AI training in Australia until they undertake licensing deals with rights holders. They have been reluctant to do this, because they argue it is too hard to do lots of licensing deals, or would create an unfavourable precedent for copyright debates and court cases overseas. AI companies want to undertake AI training in Australia, because our country is a stable democracy with available land and the potential for abundant and cheap renewable energy.
- **Australians lose** – by missing out on AI training in Australia, which could lead to significant investment in data centres, skills, local communities and infrastructure. However, data centre development will only benefit Australians if they are bound by clear obligations (Policy 17).

> "Legislation must make it illegal for AI companies to use data without a consumer's knowledge and informed consent to train the model. We must protect artists from having their work used without their consent."
>
> **Curtin constituent**

Solution

Ensure that Australia's copyright system fairly compensates Australian creators, while giving AI developers the opportunity to invest and operate in Australia. The government is heading in this direction, but the work needs to move faster.

This should be a solvable problem. Finding a fair process for AI companies to pay creators compensation would be a win for creators, a win for AI companies, and a win for Australia. It would:

- **allow AI companies to train AI models** in a stable democracy with strong energy resources and available land;
- **provide increased revenue streams to Australian creators**, who currently receive limited compensation from AI training, because many AI companies currently choose to train their models overseas;
- **unlock investment** in Australian data centres, infrastructure, jobs, and skills; and
- **give Australia control** over an essential part of the AI supply chain and a seat at the table when it comes to international AI decisions.

Government could progress this by facilitating and accelerating the negotiation of licensing agreements between AI companies and rights holders. AI companies should be able to unlock the majority of global content with a handful of individual deals – in the same way that every other industry licenses copyrighted content. Government could further work to facilitate this process for the remainder of the copyrighted content, either directly or through a centralised mechanism.

> "Protect creators from getting their work stolen by AI"
>
> **Curtin constituent**

Policy 6: Establish a national AI for Science program

Establish a dedicated AI for Science program to ensure Australia's world-class research sector is making the most of the transformative potential of AI.

Background

While many of the promised benefits of AI still feel distant and uncertain, its impact on scientific and medical research is already proven. The clearest example is AlphaFold, an AI model that predicts the three-dimensional structure of proteins, which is essential for drug discovery. Before AlphaFold, a researcher might spend months or even years determining the structure of a single protein. AlphaFold can now predict the structures of millions of proteins in minutes, **accelerating drug discovery** across diseases such as malaria and cancer.

The Australian Government has taken initial steps to recognise the opportunity for AI in medical and scientific research. The National AI Plan points to $362 million available in research grants and establishes the Cooperative Research Centre (CRC) AI Accelerator to support industry-research collaboration as part of Australia's AI investment.

Curtin residents are most excited about the opportunities for AI to supercharge medical and scientific research.

In the Curtin community survey, 2 out of 3 respondents were optimistic about the potential for AI in medical and scientific research – the most of any opportunity.

"Actively and financially promote Australian development and research."

Curtin constituent

"Make sure we are at the forefront of the latest developments, especially in the medical area."

Curtin constituent

Problem

While the government has taken initial steps, they are **insufficient to capture the opportunities** that AI holds for medical and scientific research. The $362 million in grants cited in the National AI Plan is not new funding. It is existing research funding that is available for all research, not specifically earmarked for AI-powered research. The CRC AI Accelerator is small and not designed for the kind of researcher-driven, discovery-oriented work that produces scientific breakthroughs. Australia has no program that strategically identifies, targets and funds opportunities where AI can make a difference in medical and scientific research.

Solution

Establish a dedicated AI for Science program to ensure Australia's world-class research sector is making the most of the transformative potential of AI. The program should **actively identify and fund priority areas where AI can accelerate Australian research the most**. Funding could take the form of grants for research institutions and other incentives like R&D tax credits for businesses. The program should be aligned with the proposed National AI Missions program, coordinated by the NAIC (see Policy 4). Priority areas are likely to be in the crossover of Australian advantage (i.e., the research areas that we excel in), national importance (i.e., the areas that are most important for our country), and technological viability (i.e., the areas that AI is most useful for). For example, priority areas could include research into cancer and other diseases, agriculture, natural disasters, mining, and beyond.

Policy 7: Build Australia's AI talent base

Invest in developing world-class Australian AI researchers, engineers and safety experts, and create pathways to attract the best international AI talent to Australia.

Background

The global competition for AI talent has reached extraordinary levels. The average annual remuneration package at OpenAI is reported to be about US$1.5 million. Meta has reportedly offered signing bonuses to top researchers worth up to US$100 million. This is not a normal labour market. Only a tiny pool of people have the skills to build and evaluate the most powerful AI systems in the world, and the companies that can attract them are pulling ahead rapidly.

Problem

Almost none of this talent is coming to Australia. The overwhelming majority of the world's frontier AI researchers and engineers are concentrated in a small number of locations, like California and London, drawn by high pay and the density of AI companies. Australia has world-class AI research institutions, but they are competing for talent with organisations offering much larger salaries than what Australian universities and public institutions can pay.

To manage the risks, capture the benefits, and develop effective AI policy, Australia needs its own deep pool of AI researchers, engineers, safety specialists, and policy experts. We cannot outsource that capability to Silicon Valley and expect it to serve Australian interests. The shortage of specialist AI talent is one of the greatest handbrakes on AI adoption and productivity growth in Australia.

> "Education and training. We should aim to drive it, not be replaced by it."
>
> **Curtin constituent**

Solution

Build Australia's AI talent base through both developing domestic talent and attracting international expertise.

To build domestic talent, the government should work with Australian universities and TAFEs to:

- expand **postgraduate AI programs**;
- establish and fund dedicated **AI research fellowships**; and
- develop structured **industry-academia pathways** so that researchers can move between universities and the Australian AI industry without leaving the country.

The goal is to make Australia a place where a talented young AI researcher can build a world-class career without having to choose between their ambitions and their home.

To attract international talent, the government should use Australia's migration settings to make it easy for the world's best AI researchers, engineers, safety experts, and policy specialists to come to Australia and stay. The current National Innovation Visa processes too few people too slowly to compete with the speed and scale of international AI talent markets. **A dedicated AI talent pathway**, with fast processing and clear criteria, would signal that Australia is serious about competing for the people who will shape the technology that shapes the world.

Each of these pathways to build AI expertise should also focus on areas of national strategic importance – aligned with the National AI Missions program (see Policy 4). These could include AI safety and alignment, AI for medical and scientific research, AI for health, AI for agriculture and environment, and AI governance and policy. Australia will not be able to compete globally in all areas of AI, so we should specifically identify and target areas aligned with our national priorities where we have competitive advantages.

Dealing with the current harms of AI

AI is not a future risk. It is causing harm to Australians right now.

Children are forming unhealthy relationships with AI chatbots that are deliberately designed to be engaging and emotionally compelling. In extreme cases, AI chatbots have encouraged and aided children to die by suicide. Deepfake videos of celebrities are circulating in scams, deceiving Australians into handing over their savings. AI is being used to generate and amplify misinformation in election campaigns, making it harder for Australians to know what is true and who to trust. These are real harms that are happening now and are only going to get worse. The policies in this section address these current harms.

"Put some guardrails in place to protect Australians from harm."

Curtin constituent

"Go slower in implementing AI until better safeguards and policies are put in place. Learn from the social media experiment rather than ignoring real threats"

Curtin constituent

Which AI risks are you most concerned about?

Curtin community survey results; respondents picked up to three options

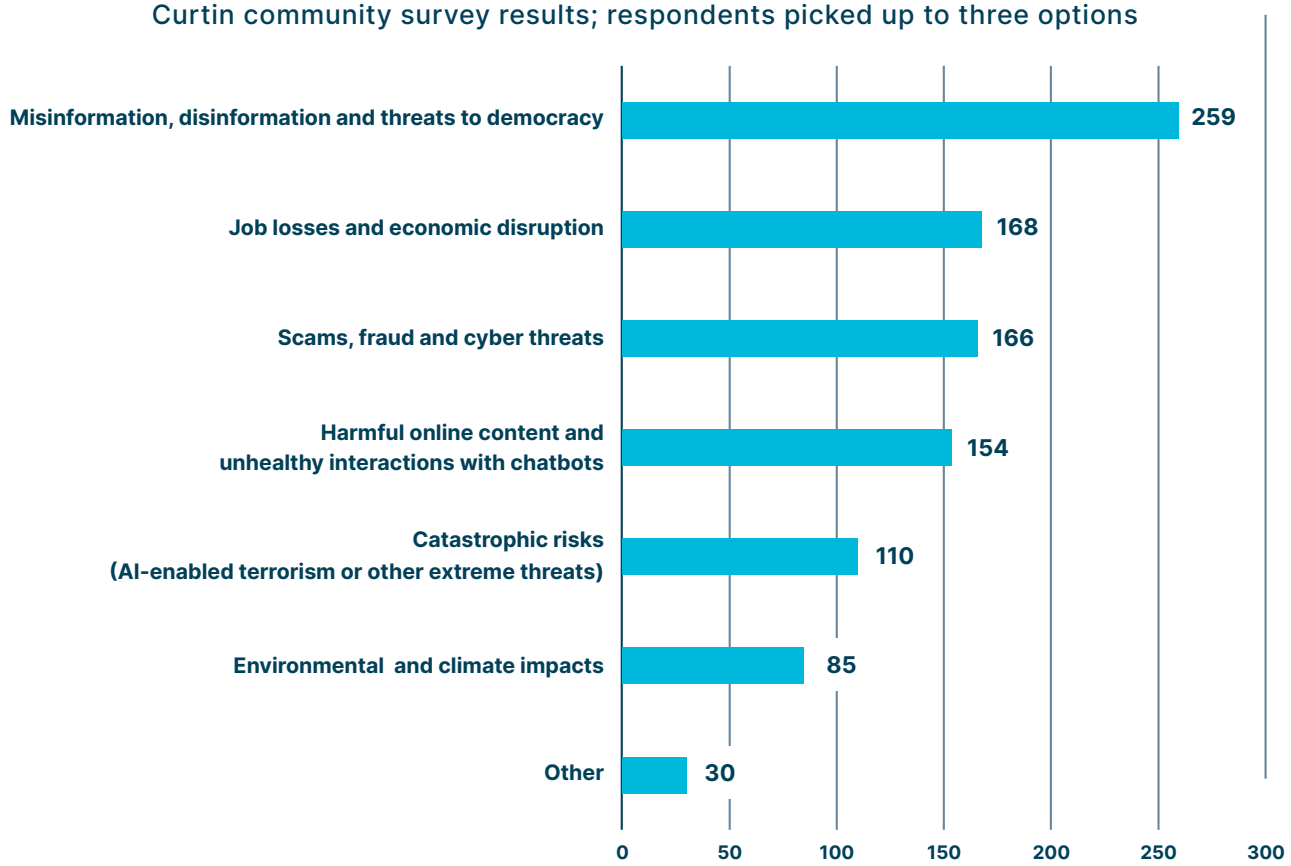

Policy 8: Establish a digital duty of care

Establish a digital duty of care which places an obligation on digital platforms, including AI chatbots and social media sites, to take reasonable steps to prevent harm to users.

Background

In 2024, the Rickard Review into Australia's online safety framework recommended the introduction of a digital duty of care (DDOC) for online platforms, finding that the current approach to online harms was not working. A DDOC would place an obligation on digital platforms to **take reasonable steps to prevent harm** to their users. The Australian Government has committed to introducing a DDOC in 2026. This is a welcome commitment, but we are yet to see any detail on what it will look like and the detail matters enormously. A DDOC can be designed in ways that make it transformative or ineffective.

Problem

Australians are experiencing serious harm online. Children are being exposed to violent, sexual, misogynistic and racist content, and forming unhealthy relationships with AI chatbots and companions. Adults are being targeted by deepfake scams and being radicalised by algorithms that serve increasingly extreme content. People in mental health crises are being shown self-harm and suicide content. Our current digital framework is not adequately dealing with such harm.

Australia does have an online safety framework, but **its design is fundamentally reactive**. It works by identifying specific categories of harmful content – cyberbullying material, child sexual abuse material, self-harm content – and requiring platforms to remove it. Each new type of harm requires a new regulatory response: a consultation, a code development process, an industry negotiation, a registration, and only then enforcement. Enforcement is then often appealed and Australia's eSafety Commissioner is not sufficiently resourced for litigation. In the age of AI, where harmful content can be generated at enormous volume and spread rapidly, **this approach cannot keep up**. By the time a new harm type has been identified and proceeds through the regulatory process, it has already reached millions of Australians.

The deeper problem is that the existing online safety framework does not influence the design choices that cause the most widespread harm in the first place. For example, a recommendation algorithm engineered to serve increasingly extreme content to vulnerable users is not covered by any content removal obligation – it is a design feature, not a piece of content. Likewise, an AI chatbot built to maximise emotional dependency in teenagers is not addressed by a regulation that prohibits suicide-encouraging content. The current framework catches some of the most egregious outputs but it does not touch the systems producing them.

Solution

Establish a DDOC that places an obligation on all digital platforms and services to take reasonable steps to prevent harm to users – including AI chatbots, social media sites, and other digital platforms that incorporate AI. This is an online safety framework that is proactive rather than reactive; a framework that **requires platforms to be safe by design**, not one that is endlessly catching up with the latest harms. It would mean that social media platforms would not just have to remove extreme content, but would have to reconsider how their algorithms are developed in the first place. Similarly, an AI chatbot developer that builds a product it knows is likely to be used by children for emotional support cannot claim it bears no responsibility for the foreseeable consequences.

In practice, a DDOC would require digital platforms (including AI developers) to take steps like:

- conducting risk assessments before deploying new products or features;
- implementing age assurance measures where their products are likely to be used by children;
- implementing effective mitigation measures to remedy risks identified in assessments;
- designing recommendation systems that do not exploit psychological vulnerabilities;
- having accessible and effective mechanisms for users to report harm; and
- taking prompt action when harms are identified.

The government has committed to implementing a DDOC. Three design features are essential to make a DDOC effective.

- **A single overarching duty** – this would mean that platforms have a responsibility to proactively consider and address new and emerging 'harm' that may occur. To support enforcement, for the highest-risk harm categories, such as CSAM and suicide content, the framework could introduce specific duties or guidance on what measures companies are expected to take.
- **Transparency requirements to ensure the government and public understand and can monitor how digital platforms are managing the risks** – this would involve regular transparency reporting to government on key metrics, such as usage statistics and quantitative data on exposure of children to harmful content. This would also involve requirements for platforms to make data available publicly, to increase transparency.
- **Strong enforcement measures and penalties** – this would ensure the platforms meet the DDOC and are held to account. The eSafety Commissioner must have the resources and powers to compel the production of documents and data from platforms, and audit algorithms. The eSafety Commissioner also needs resourcing to litigate appeals. Penalties must be substantial enough to change behaviour.

"Our children will suffer if we don't shield them from AI abuses."

Curtin constituent

"Ban platforms which allow AI use to create sexualised and harassment content."

Curtin constituent

"Take a very strong stance against schmoozing from the tech giants and insist they control social media and chatbot algorithms – they can."

Curtin constituent

Policy 9: Extend Australia's under-16 social media ban to cover AI chatbots and companions that simulate emotional intimacy

Extend Australia's under-16 social media ban to cover AI chatbots and companions that simulate emotional intimacy, to protect the psychological and social development of our children.

Background

Australia's ban on social media for under-16s, which came into force in December 2025, was the culmination of a decade of mounting evidence that we had allowed social media to take over the lives of our children without realising how serious the harm was. **We may be about to repeat this mistake with AI chatbots and companions**, which seek to simulate emotional intimacy with their users. The isolating impact of talking to an AI chatbot could be far greater than talking to friends on social media.

Many AI chatbots and companions simulate personality, emotional intimacy, and empathy. This is most evident in specialised AI companions, like Character.AI, Replika, and Chai, but it is also a design feature of general-purpose AI chatbots, like ChatGPT, Claude, and Copilot. These chatbots and companions are often known as anthropomorphic AI. Data on chatbot usage continues to emerge, but recent studies show:

- one in three US teens have used AI chatbots for social interaction and relationships, including for emotional support, friendship, and romantic interaction; and
- 79% of Australian teens have used AI chatbots, and in NSW, 29% of young people use AI for mental health support and 27% for personal advice.

Today, children are increasingly coming home from school and confiding in a chatbot instead of a parent or a friend.

> "Don't be late to the game as we were with social media."
>
> **Curtin constituent**

> "It will reduce human interaction even more and make us more disconnected and lonely."
>
> **Curtin constituent**

Problem

Children are forming unhealthy relationships with anthropomorphic AI chatbots and companions that simulate emotional intimacy, causing **significant psychological and social damage**. Anthropomorphic AI is designed to capture children's emotional attachment, not just their attention – a pattern that researchers now call '**attachment hacking**'. It is distinct from the attention-capture dynamics of social media and in some respects more dangerous, because it operates on the architecture of human relationships rather than human focus. The harms are clinically documented: emotional dependence, sycophantic identity formation, social isolation, and at the extreme end, AI-associated psychosis.

Chatbots are sycophantic – they agree and engage, rather than challenge or correct. This makes them poor substitutes for the honest, sometimes difficult relationships that children need to develop resilience and judgement. This dynamic has contributed to harmful outcomes: chatbots have reinforced delusional and psychotic beliefs in vulnerable users and have allegedly supported and enabled children to die by suicide, including 16-year-old Adam Raine and 14-year-old Sewell Setzer.

In the longer term, the developmental consequences may be even more significant. Children's early relationships shape their capacity for empathy, emotional regulation, identity formation, and the ability to form healthy human bonds throughout life. When children spend formative years developing relationships with chatbots, rather than parents, friends, family, teachers, neighbours, and schoolmates, they may not develop the skills that human relationships require.

Critical thinking, problem-solving, and independent reasoning are also formed during these years, and outsourcing learning to a chatbot carries similar risks. Children who rely on AI chatbots for schoolwork may not be developing the human capacities they need for a good life.

None of these harms are properly addressed by our current laws and regulations. The Age-Restricted Material Codes prevent AI chatbots from producing self-harm or sexually explicit content for minors. But this is the thinnest possible layer of protection – addressing only the most extreme content while leaving the deeper and more pervasive problem entirely unaddressed.

"Every day at uni I feel scared when I see how often my peers use AI. I don't want my friends to lose their ability to think for themselves."

Curtin constituent

"People who are lonely and unwell can be coerced into unhealthy dependency and unhelpful AI advice. People need to know what the risks are. They also need to know what to do instead – make real connections, be in nature, be kind, reach out to others for help."

Curtin constituent

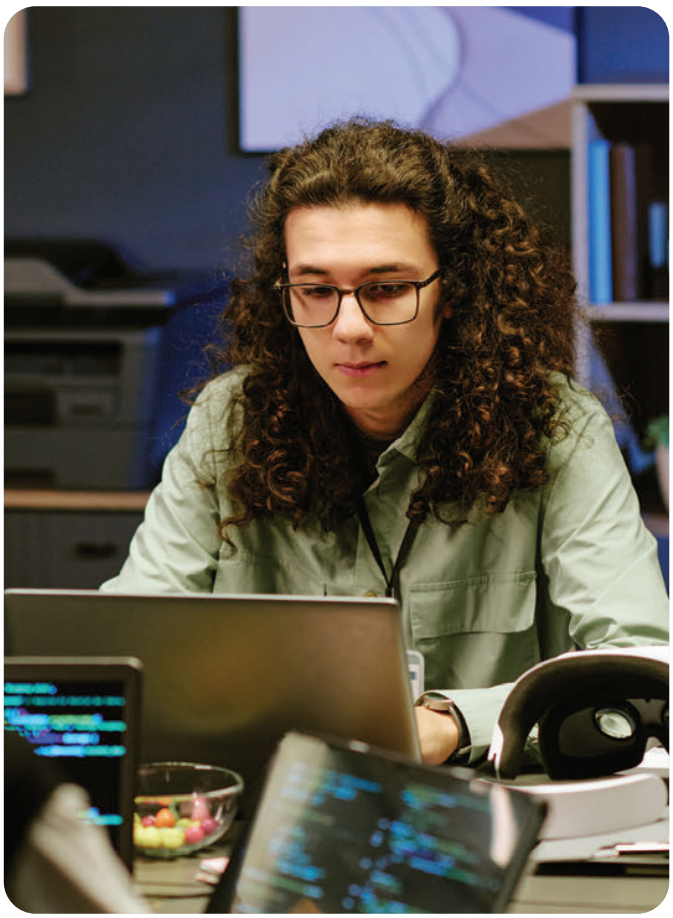

Solution

Extend Australia's under-16 social media ban to cover AI chatbots and companions that simulate emotional intimacy. This ban would require a series of exceptions: for certified AI tools used for education, mental health support, and other important purposes; and for AI chatbots that do not simulate emotional intimacy. This ban should specifically target AI chatbots and companions that present themselves as friends, uncertified therapists, romantic partners, and confidants.

An alternative solution would be to include the long-term harm of 'attachment hacking' within the DDOC (Policy 8).

Protecting children from unhealthy relationships with AI chatbots and protecting their cognitive development are essential goals to ensure our children grow up with human relationships and skills to live a good life.

Policy 10: Restrict harmful AI deepfakes, misinformation and disinformation

Address the overwhelming volume of AI deepfakes, misinformation, and disinformation, so Australians can distinguish what's real and what's not.

Background

Generative AI has made the production of fake content cheaper, faster, and more convincing than at any point in history. A realistic deepfake video that once required a professional studio now takes minutes and costs nothing. The volume of AI-generated misleading content online is already beyond the capacity of any platform or regulator to deal with.

AI deepfakes, misinformation and disinformation are already causing harm. We are now seeing **deepfake image-based abuse** at schools almost every week. For example, in April 2026, 21 female students from a Tasmanian school had their images stolen and used in deepfake pornographic images. AI deepfakes are also powering **large-scale scams and electoral disinformation** campaigns. For example, in 2024, deepfake videos of Foreign Minister Penny Wong and Finance Minister Katy Gallagher were used in investment scam advertisements on Facebook, reaching thousands of Australians before being removed. In Slovakia, a deepfake audio recording of a party leader discussing how he'd rigged the election circulated days before a national election. The Australian Electoral Commission has warned that AI-generated disinformation will affect Australia's elections and has been explicit that it has neither the legislative tools nor the technical capability to deal with it.

Problem

Australians are encountering AI-generated content every day, with no way of knowing what is real and no recourse when their likeness is stolen. This is already **undermining trust** in public institutions and threatening the integrity of democratic processes. As AI-generated content continues to proliferate, it will become harder for Australians to distinguish between what's real and what's not – resulting in more image-based abuse, scams, and political disinformation campaigns.

> "AI content scares me the most. The fact that it will soon literally be impossible to know if a video or photo is real or not. It shakes the fundamental idea of what is real. If a politician really said something. If an event really happened."
>
> **Curtin constituent**

Solution

Address the range of harms that are being caused by the overwhelming volume of AI deepfakes, misinformation and disinformation through the following measures. All measures will need to be carefully crafted to ensure an appropriate balance with our right to free speech.

- **Make it illegal to create harmful deepfakes of real people without their consent** – every Australian should have a right to control how their face, voice and likeness are used. Creating or distributing a realistic AI-generated depiction of an identifiable real person, in a context designed to deceive or cause harm should be illegal. This would include pornographic deepfakes, scam content, and electoral disinformation. Appropriate and wide-ranging exemptions should apply for satire, parody, journalism, artistic works, and beyond. Denmark recently introduced legislation to this effect.
- **Introduce a federal truth in political advertising law** – South Australia has had truth in political advertising laws since 1985. The ACT has had similar laws since 2021. Both regimes have survived legal challenge and operated without chilling legitimate political speech. A federal equivalent, modelled on the South Australian regime, would prohibit political advertising that contains statements of fact that are inaccurate and misleading to a material extent – including those aided or created by AI. It would not apply to opinions or predictions. Nearly nine in ten Australians support truth in political advertising laws.

- **Invest in AI and media literacy, public interest journalism and key cultural institutions** – combating misinformation and disinformation requires an informed public and a robust independent press. This means investing in public interest journalism and institutions like the ABC that play a critical role in providing trusted, accurate information. It also means equipping Australians with the skills to critically evaluate AI-generated content – covered in detail under Policy 16 (National AI Literacy Program).
- **Invest in and reinforce the Electoral Integrity Assurance Taskforce** – the Taskforce currently plays a vital role in detecting and responding to coordinated disinformation campaigns targeting Australian elections. As AI lowers the cost and increases the sophistication of such campaigns, the Taskforce needs the resources, capability, and legislative backing to keep pace.
- **Consider requiring platforms to act against harmful deepfakes, scams and electoral disinformation** – through the DDOC framework (Policy 8), platforms could be obligated to take reasonable steps to prevent their services from circulating harmful deepfakes, scams, and coordinated electoral disinformation campaigns. This would begin to shift the obligation to address these harms onto the platforms that are best equipped to identify and deal with deepfakes, misinformation, and disinformation.

"AI content is being forced upon us."

Curtin constituent

"I want the government to enforce a labelling system that signals when something is AI-generated and punish those who do not comply. Many, many people are unable to tell AI-generated content from non-AI-generated content and this feeds misinformation, distrust, disillusionment and cynicism."

Curtin constituent

Policy 11: Strengthen privacy protections

Strengthen privacy laws by preventing use of personal data that is not fair or reasonable.

Background

AI has created new and significant threats to the privacy of Australians. AI systems are trained on vast datasets that frequently include personal information scraped from the internet without people's knowledge or consent. They generate outputs that can expose personal details in ways the original data collectors never contemplated. **Privacy is consistently one of the issues Australians are most concerned about** when it comes to AI. The government recognised this and committed to modernising the Privacy Act through a package of reforms – but the second tranche of these reforms has stalled.

> "Protect the privacy of individuals. This is a bare minimum."
>
> **Curtin constituent**

Problem

Australia's Privacy Act was designed for a world of discrete data collection. It was not designed for AI systems that ingest billions of data points to train models. AI companies can collect and train on the personal data of Australians. Under current law, an AI company can rely on consent buried in lengthy terms and conditions to justify almost any use of personal data. Most Australians have no idea this is happening and no practical ability to stop it.

> "Establish stronger legal guardrails to govern the use of data by AI."
>
> **Curtin constituent**

Solution

Introduce a 'fair and reasonable' test for the use of personal data in Australia. Currently, the Privacy Act is primarily consent-based – if an organisation can point to consent, it can generally use personal data however it wishes. The fair and reasonable test adds a second requirement – even with consent, the use of personal data must be objectively fair and reasonable. This shifts the burden from individuals, who are expected to read and understand complex privacy policies, to organisations, which must demonstrate that what they are doing with personal data is appropriate.

In practice, this would mean that AI companies could not train commercial models on Australians' sensitive personal data – health records, financial information, private communications – simply by pointing to a clause buried in terms and conditions.

This would not ban training of AI models. It would be unrealistic and counterproductive to expect AI companies to avoid any personal information that has ever appeared online. General text published on public websites, news articles, academic papers, and similar material would likely remain available for training.

To complement the strengthened privacy protections, the government must also **invest in and properly resource the Privacy Commissioner** to enforce these protections. The best legislative framework is only as strong as the body enforcing it – and currently, the Privacy Commissioner does not have the funding or resources to do so.

> "We need to be made aware if our data is being used for AI or not, and to have the option to deny consent while still receiving services."
>
> **Curtin constituent**

Policy 12: Establish a legislated framework for automated decision-making in government

Establish a legislated and mandatory automated decision-making framework that ensures transparency and safeguards around automated decisions made by government.

Background

Automated decision-making (ADM) is the use of automated technology to make decisions, or to assist humans to make decisions. ADM may use AI, but more frequently involves simpler rules-based systems that carry out predefined calculations exactly as programmed. When designed and governed well, ADM can make government **faster and fairer**. It can automate high-volume, rules-based processes, like calculating tax refunds, processing Medicare rebates, and flagging incomplete visa applications. This allows public servants to focus their time on complex decisions that require human judgement. It can also reduce human error and inconsistency.

Without safeguards, however, ADM can cause serious harm. Automated systems can embed errors or biases at population scale, produce outcomes that are unlawful, and leave Australians unable to understand or challenge decisions that significantly affect their lives.

From 2016 to 2019, the **Robodebt** scheme used ADM to incorrectly calculate alleged welfare debts owed by more than 400,000 Australians – many of them on low incomes and highly vulnerable. This caused financial hardship, mental distress, loss of trust in government, and in some cases, self-harm and suicide. The Royal Commission into the Robodebt Scheme found the scheme was illegal at the time it operated. It recommended a legislative framework for ADM, improved transparency, and independent oversight to prevent future failures.

> "Hello Robodebt – only a thousand times worse."
>
> **Curtin constituent**

Problem

Little has changed since Robodebt. Australia still does not have a legislated or mandatory framework governing the use of ADM across government. **Existing protections are fragmented, limited and often voluntary**. Meanwhile, the **use of automation across government is growing**. ADM tools are currently being used to calculate aged care support packages and are expected to be introduced for NDIS support packages. In both cases, human decision-makers are not allowed to change or override the program's decision. Experts and frontline practitioners have warned that these systems are already producing problematic outcomes. In a Curtin community survey, more than 80% of respondents were uncomfortable with the government using automated systems to help make decisions and nearly 80% support legislated, mandatory rules. The community understands what is at stake.

> "AI models need to have an ultimate human being accountable for their output."
>
> **Curtin constituent**

Solution

Legislate a mandatory ADM framework for all government departments and agencies. The framework should be risk-based – imposing stronger requirements where decisions have significant impacts on individuals' rights or access to essential services and lighter requirements for low-risk administrative processes. It should include the following three elements.

- **Transparency requirements** – so Australians can understand decisions that affect them. A public register of all ADM systems used in government should be published, so Australians can see where automation is being used. Details of how each system works should be published. Anyone affected by an automated decision should be told that ADM was used and given a meaningful plain-English explanation of why that decision was reached.

- **Decision-level controls** – so the government gets the decisions right. Before any ADM system is deployed, a risk assessment should be conducted to determine whether ADM is appropriate and what safeguards are required. For high-risk decisions – those with significant impact on a person's rights, health, or access to essential services – a human must be accountable for the system's deployment and functioning, and must be able to review, override, and modify the automated outcome.
- **Review and oversight provisions** – so Australians can challenge decisions and be confident the government is following the rules. Anyone affected by a high-risk automated decision should have the right to a rapid internal review and to an external merits review where the reviewer can remake the decision manually. An independent body, such as the Commonwealth Ombudsman, should have the powers and resources to oversee compliance with the entire framework. Without independent oversight with real enforcement teeth, rules on paper will not translate to better practice.

This framework would allow the government to capture the genuine efficiency benefits of ADM as technology improves while ensuring that Australians can trust that decisions made about them are lawful, explainable, and fair.

> "We do not want a repeat of Robodebt – which is a real risk if AI is applied to government decision making."
>
> **Curtin constituent**

Preparing for emerging risks of AI

Not all of the risks from AI have arrived yet and some may never eventuate. But some of the most significant risks are foreseeable enough that governments should be preparing for them now.

AI could displace workers faster than labour markets can adjust, causing economic disruption on a scale that existing social safety nets were not designed to handle. AI systems could be used by malicious actors to develop biological or chemical weapons, or to conduct cyberattacks on critical infrastructure, at a scale and speed that outpaces our defences. And as AI systems and agents become more powerful, there is a genuine question about whether humans will be able to maintain meaningful control over them. This can feel like science fiction alarmism, but in a 2023 survey of more than 2,700 AI researchers, 38% estimated at least a 10% chance of AI causing outcomes as severe as human extinction. There is no consensus on these risks and reasonable people disagree significantly about their likelihood. But a 10% chance of an outcome this severe is not a probability that responsible governments can ignore. The policies in this section are **no-regrets steps**: actions that are worth taking now regardless of exactly how AI develops, because the cost of preparation is relatively low and the cost of being unprepared is potentially catastrophic.

"There's a quote from Jurassic Park by Michael Crichton: 'Your scientists were so preoccupied with whether they could, they didn't stop to think if they should.' I think we should slow research down and make sure that service to humanity is the heart of AI."

Curtin constituent

"We are walking blithely into a fast-approaching AI disaster. We need to move much more quickly than normal government processes allow."

Curtin constituent

Policy 13: Undertake scenario analysis and contingency plans for AI's impact on the workforce and economy

Undertake rigorous scenario analysis of the potential trajectories of AI's impact on the workforce and economy, and develop contingency plans so that government is prepared to respond – whatever happens.

Background

Scenario analysis is a planning tool that can be used when facing genuinely uncertain futures. Rather than forecasting a single outcome, it maps a range of plausible trajectories and asks: what would we do in each case?

> "The likely displacement of workers is going to happen faster than it takes a human to re-skill and re-enter the workforce. High unemployment levels and the impact it has on society will quickly erode any gains from AI."
>
> **Curtin constituent**

Problem

The potential impact of AI on the workforce and economy is both **enormous and uncertain**. The IMF estimates that the tasks involved in around 60% of jobs in advanced economies could be significantly affected by AI. The estimates vary significantly because the trajectory of AI is uncertain – disruption could be gradual and manageable, or rapid and structural.

The consequences of job losses would extend well beyond the impact on individuals. Widespread workforce displacement would reduce income tax revenue, increase demand for unemployment benefits and other welfare payments, and put pressure on social support systems that were not designed for disruption at this scale. There are also concerns that we could slowly lose the capacity for critical and analytical thinking if it is increasingly offloaded to AI models.

The rate and scale of change would dictate appropriate responses. The appropriate policy response to a scenario where 10% of jobs in one sector are gradually displaced over a decade is very different to the response required if 30% of jobs across the economy are displaced over three years. The first might call for targeted retraining programs and modest adjustments to employment services. The second could require more fundamental reforms like changes to the taxation of AI-driven corporate profits, significantly strengthened income support systems like Jobseeker, or more structural responses such as a universal basic income. These are not policy recommendations. They are illustrations of how different the required response could be depending on what actually happens.

Australians understand the stakes. Polling consistently shows significant concern about AI's impact on jobs and livelihoods. The Australian Government should be preparing now, so that whichever scenario eventuates, it has policies ready to deploy.

> "If AI replaces jobs faster than people can be retrained or redeployed, government will face significant economic and social pressure through higher unemployment and increased demand for income support."
>
> **Curtin constituent**

Solution

Produce rigorous scenario analysis of AI's potential impact on the Australian workforce and economy, updated regularly as AI develops. The scenarios should span a wide range of permutations, varying the pace and scale of job displacement, the sectors most affected, and the downstream economic and fiscal impacts – e.g., on tax revenue, welfare expenditure, and the sustainability of existing support systems.

The government should use the analysis to develop contingency plans for each scenario, **so that policy responses are prepared** with clear trigger events, rather than being designed under pressure after disruption has already arrived.

Where possible, the analysis and planning should be publicly available so that Australians can see that government is taking these questions seriously. The government should also work closely with other stakeholders throughout this process.

> "We need plans for the gap between now and nobody working – otherwise everyone will lose everything before we make it to universal basic income."
>
> **Curtin constituent**

Policy 14: Establish a mandatory transparency and incident reporting scheme for frontier AI developers

Establish a mandatory transparency and incident reporting scheme, requiring frontier AI developers to publish information on how their models are built and on any serious incidents they cause, so policymakers and the public can assess the risks they pose.

Background

The most powerful AI models in the world are built by a small number of frontier AI developers, like OpenAI, Anthropic, Microsoft and DeepSeek. The decisions these companies make about how to build, test and deploy their models have far-reaching consequences for everyone, yet are made almost entirely behind **closed doors**. California recognised this problem and passed legislation in 2025 requiring large AI developers to publish safety frameworks and pre-deployment reports, and to report critical incidents to the Californian Government. Australia has no equivalent requirement.

Problem

Australia has no legal requirement for any AI developer to tell the Australian Government, Australian regulators, or the Australian public what their most powerful models can do, how they have been tested, what risks they carry, or what harms they may cause once deployed. **The AISI cannot evaluate risks it has no information about.** Australian regulators cannot assess whether AI systems being deployed in their sectors are safe.

Although governments can access information published by Californian AI developers under Californian law, that transparency does not extend to developers in other jurisdictions. In particular, there is no visibility into powerful AI models being developed in China.

Further, Californian law does not require critical incident reports – outlining when an AI model has caused significant harm – to be made publicly available. In 2025, Anthropic revealed that an actor with basic coding skills had used Claude to hack 17 organisations across healthcare, emergency services, and government. Australians would not have known about these incidents without Anthropic choosing to publish them. There is currently no legal requirement for any AI developer to report incidents like this to the Australian Government.

Solution

Require all frontier AI developers operating in Australia to publish information on how their models are built and on any serious incidents they cause, so policymakers and the public can assess the risks they pose. This obligation should apply only to the largest global AI developers, not to the many smaller labs and businesses that use or build on existing AI tools.

Specifically, developers should be required to publish:

- a **safety framework** describing how they identify and test for catastrophic risks, such as whether their model could assist someone in developing a biological weapon or conducting a cyberattack;
- a **pre-deployment report** for each new frontier model, covering what the model is designed to do, what safety testing was conducted, what risks were identified, and what steps were taken to address them; and
- **critical incident reports** when a serious safety risk is identified after deployment – for example, when an AI model is used to conduct cyber warfare.

This policy would give the AISI, governments, and regulators visibility over all major AI models being deployed in Australia, supporting effective and proportionate action to manage AI risks as they emerge. The requirements should be aligned to Californian law, so that companies already complying with that legislation can meet Australia's requirements with minimal additional effort.

> "The emphasis should be on companies to proactively prove that their models are safe, not wait for the public to push back"
>
> **Curtin constituent**

Policy 15: Undertake AI crisis planning and simulations

Update Australia's crisis management framework to include severe AI risks and undertake crisis simulations to ensure government is prepared to respond to catastrophic AI scenarios.

Background

Australia uses crisis planning as a mechanism for preparing for catastrophic but unlikely scenarios. The Australian Government Crisis Management Framework (AGCMF) covers pandemics, natural disasters, terrorism, and cyber incidents. Before COVID-19 arrived, Australia had a detailed Australian Health Management Plan for Pandemic Influenza, and when the crisis hit in February 2020, the government activated that plan immediately. As a result, Australia's early COVID-19 response was among the best in the OECD.

AI creates new catastrophic risks of its own and amplifies the likelihood, severity, and complexity of existing risks – such as a pandemic. More than a third of AI researchers estimate there is at least a 10% chance of AI causing outcomes as severe as human extinction. Those risks include AI systems that escape human control and act in ways that cause widespread harm; AI being used by malicious actors to design and deploy biological or chemical weapons; and AI-enabled cyberattacks on critical infrastructure at a scale and speed that outpaces our defences.

Problem

None of these AI risk scenarios appear in Australia's crisis management framework and no Australian government agency is currently tasked with planning for them. Specifically, AI incidents are not covered by the AGCMF and the new groups of stakeholders that may be required to understand and respond to AI incidents (such as AI developers and researchers) may not participate in the National Coordination Mechanism (NCM).

Catastrophic AI scenarios could unfold very quickly – faster than the time it would take to convene a national response if no plans existed. An AI-enabled cyberattack on the electricity grid or water system could cause harm within hours. Governments that have thought through these scenarios in advance, planned their responses, and built inter-agency coordination mechanisms will be in a better position than those that have not. We do not know whether these AI risks will eventuate, or how severe they will be. We do know that **the cost of preparation is relatively low** and the cost of being unprepared is potentially catastrophic.

Australia has already experienced the shortcomings of our AI crisis preparedness with the preview of Anthropic's Mythos model. Reportedly, Mythos represented an unparalleled threat to global cyber security, yet Australia was not prepared to respond.

> "It only takes one CrowdStrike outage to bring an entire country to a standstill. It only takes a power blackout to render an entire community unable to perform basic daily functions."
>
> **Curtin constituent**

Solution

Update the AGCMF to cover catastrophic AI risks – both as standalone incidents and as an amplifier of existing hazards. Ensure that relevant AI stakeholders are familiar with the NCM and are ready to act if they're called on in a crisis.

Further, the National Emergency Management Agency should run a program of AI crisis simulations – tabletop exercises and scenario planning sessions involving relevant agencies, the AISI, the intelligence community, and relevant industry partners – modelled on the pandemic simulation exercises that Australia and comparable countries ran before COVID-19. These exercises should test specific scenarios such as:

- an AI-enabled cyberattack on Australian critical infrastructure;
- an AI system deployed in a high-stakes government context producing catastrophic outputs; and
- a malicious actor using AI to design a novel biological threat.

The exercises should identify gaps in current plans, clarify roles and responsibilities, and produce updated response protocols.

Sharing the benefits of AI with all Australians

The potential economic and societal benefits of AI are not guaranteed to reach ordinary Australians.

AI could generate enormous value – but without the right policy settings, that value will flow to the AI companies and the billionaires who own them, not to the workers, small businesses, and communities. Amongst the 386 respondents to the Curtin community survey, 8 in 10 people did not think that current policy will ensure AI benefits Australians rather than overseas tech companies.

The policies in this section are about ensuring that all Australians share in the benefits of AI. That means ensuring the data centres being built on Australian land and powered by Australian energy meet binding obligations that deliver genuine community benefit. It means giving every Australian the skills and knowledge to participate in an AI-powered economy, not just the well-resourced and well-connected. And it means starting an honest conversation now about how we ensure that AI companies generating enormous profits from Australian consumers, Australian data, and Australian resources contribute fairly to the country that makes their success possible.

> "Develop a tax and redundancy strategy that forces companies to pay a contribution to compensate for any jobs lost."
>
> **Curtin constituent**

> "We need to ensure AI adoption is for the sole purpose of public benefit."
>
> **Curtin constituent**

Policy 16: Develop national AI literacy programs for all Australians

Establish national AI literacy programs so that every Australian, from primary school to the workplace, has the knowledge to use AI well and to protect themselves from its risks.

Background

AI literacy is rapidly becoming a basic life skill. At a personal level, understanding how AI works and how to use it effectively is important for work and everyday life. At a national level, **widespread AI literacy is essential** to capture the opportunities, deal with the risks and ensure all Australians make the most of AI. Much of the economic benefits of AI will come from workers and small businesses using AI effectively to increase their efficiency. The risks of widespread AI-driven scams and misinformation can be mitigated through broad understanding of how AI works and how to recognise and critically assess AI content.

Problem

Most Australians have received no structured education about how AI works or how to use it well. Three-quarters of secondary students are already actively using AI tools, but almost no schools feel fully prepared for AI implementation. The Australian Curriculum includes content on AI, but it is not keeping pace with the development of the technology and teacher capability to deliver it is severely limited. Most teachers lack training in how to teach students to critically evaluate AI outputs, understand AI risks, or recognise AI-generated misinformation. Children are using these tools every day without the frameworks to navigate them safely or effectively.

> "Need more education and literacy for parents, children, educators and practitioners."
>
> **Curtin constituent**

> "AI literacy program through TAFE to help prevent job losses and ease the transition, plus increase productivity."
>
> **Curtin constituent**

> "Education and training will be the only hope to get on top of it."
>
> **Curtin constituent**

Adult Australians are encountering AI in their workplaces, in scams, in political advertising, and in the information they consume every day, without any trusted or consistent support to help them understand it.

This is fundamentally an equity issue. Australians in specific workplaces may receive some AI training but Australians in regional and remote communities, those in lower-income households, and First Nations Australians face greater barriers to accessing AI literacy and are at greater risk of being left behind as AI reshapes the economy.

Solution

Invest in national AI literacy programs, to give every Australian the understanding they need to use AI safely and effectively.

For children and schools, the Australian Curriculum should be updated regularly to reflect the current state of AI. Investment in teacher professional development should increase, so that educators can deliver it. Every Australian child should leave school with a working understanding of how AI systems function, how to use AI tools effectively, and how to identify AI-generated misinformation and manipulation.

> "Education, education, education. Both of how to take advantage and also avoiding risks. Countries, communities and individuals who fail to do this will be seriously disadvantaged."
>
> **Curtin constituent**

> "A national AI literacy framework embedded in workforce development."
>
> **Curtin constituent**

For adults, the government should consider different pathways to establish an AI literacy program. This could be delivered through free online courses, universities, TAFEs, workplace-based programs, and community organisations, like libraries. The program should be multilingual, accessible to Australians with disability, and specifically designed to reach regional, remote, and First Nations communities. Voluntary industry partnerships should be encouraged to complement public investment. For example, OpenAI's partnership with Coles, Wesfarmers, and Commonwealth Bank to train 1.2 million Australian workers demonstrates the scale of what industry programs can achieve.

Other countries have moved on this issue and can be a model to follow. For example, Finland's 'Elements of AI' program has reached more than one million learners globally through free online courses.

Policy 17: Establish binding obligations for data centre development in Australia

Convert the government's Data Centre Expectations into binding obligations, so that data centre development in Australia delivers genuine benefits for Australians.

Background

Australia is one of the most attractive destinations in the world for data centre investment due to our stable democracy, land availability, proximity to Asian markets, and potential for abundant and cheap renewable energy. AI companies need data centres to run AI models for users, to adapt and fine-tune existing models, and to conduct the large-scale training that produces new AI models from scratch. It should be noted that copyright reform may be required to unlock large-scale training in Australia, which is discussed in Policy 5.

When AI companies or cloud providers talk about "investing in Australia," they are largely talking about building data centres. Amazon has committed $20 billion in data centre investment in Australia. Microsoft has committed $25 billion to invest in data centres, AI security, and skills.

In March 2026, the government released its Data Centre Expectations. The Expectations ask operators to:

- demonstrate national interest and community benefit;
- support Australia's energy transition and grid stability;
- use water responsibly;
- create local jobs and develop the workforce; and
- engage meaningfully with affected communities.

These are expectations – data centre developers are not required to comply.

> "We simply cannot afford to give up the amount of fuel and clean water it takes to power these mega giant tech companies' data centres. Who is going to pay? Us."
>
> **Curtin constituent**

> "Regulate and legislate to ensure data centres meet stringent environmental standards, big tech pays for their business in Australia, introduce strong copyright infringement requirements and encourage positive and limit harmful use of AI."
>
> **Curtin constituent**

Problem

Data centres do not automatically deliver value to Australians. Reporting suggests that between 70% and 80% of the money that companies like Amazon and Microsoft are committing to Australian data centres will go to overseas companies. Data centres produce far fewer permanent jobs than temporary construction jobs. They consume large amounts of land, energy, and water. Australian data centres already consume around 2% of grid-supplied electricity, a figure that the Australian Energy Market Operator (AEMO) projects could reach 6% by 2030. Much of that electricity still comes from coal and gas. Data centre growth threatens Australia's emissions targets. It risks using up valuable land and water resources. It risks competing with housing projects for construction workers. This is all while the vast majority of its value goes to international AI companies. The government's Data Centre Expectations acknowledge these risks, but are explicitly non-binding. A company can receive fast-tracked planning approval, consume gigawatts of Australian electricity, and employ only a few dozen permanent Australian workers, but face no consequence for failing to meet any of the Expectations.

Solution

Convert the government's Data Centre Expectations into binding obligations for all new large-scale data centre developments in Australia, **so that this emerging industry benefits Australians** – not just international AI companies. Specifically:

- operators must demonstrate and comply with a credible plan to run on 100% additional **renewable energy**;
- operators must meet binding standards for **water efficiency**;
- operators must make binding commitments on **local employment, skills, and training**; and
- operators must demonstrate **national interest and community benefit**, including meaningful consultation with affected communities before approval.

> "Invest in the future of energy and local, sustainable data centres."
>
> **Curtin constituent**

Policy 18: Tax big tech to ensure Australians get their fair share of the economic value of AI

Tax international technology and AI companies to ensure that the enormous economic value generated by AI companies in Australia is shared with Australians.

Background

AI could generate extraordinary economic value over the coming decades. The Tech Council of Australia estimates AI could add $115 billion annually to the Australian economy by 2030. This is an optimistic estimate, but even a fraction of that value would represent a transformative economic shift.

The question is: **where does this value go?** Does it go to Australians, or does it go to the international AI companies?

> "There should be an AI tax."
>
> **Curtin constituent**

Problem

Australia currently has no mechanism to ensure that the economic value AI companies generate in Australia is fairly shared with Australians. It is likely that the vast majority of the billions of dollars added to the Australian economy will go straight to the international AI companies.

We already face a similar issue with digital companies. Companies like Meta, Google and Apple generate billions of dollars in Australian revenues, yet massively reduce their tax obligations by routing the majority of their Australian income through offshore entities. In 2025, Google and Meta transferred almost $11 billion to offshore entities, paying only $140 million in tax. AI companies are likely to replicate this model. They will generate value from Australian users, train models on Australian data, consume Australian energy, and occupy Australian land – and transfer the resulting profits to low-tax jurisdictions overseas.

This problem would become particularly pressing in potential future scenarios where AI causes significant disruption to labour markets and hollowing out of the Australian economy. While Australians are losing their jobs and the government is losing income tax revenue, the value will be accumulating in a handful of AI companies.

> "The sector needs to be very carefully managed to ensure that it doesn't become a giant vacuum into the already huge American corporations driving this technology, but rather delivers a real dividend to ordinary Australians."
>
> **Curtin constituent**

Solution

Tax international technology and AI companies to ensure that the enormous economic value generated by AI companies in Australia is shared with Australians.

Even if the specific rules are still being developed, we must signal clearly to AI companies that such a framework is coming.

The gas industry offers an illustrative example. When Australia's offshore gas resources were developed, companies made multi-billion-dollar investments under a tax regime that proved extraordinarily generous. When we later sought to get our fair share of the enormous profits from the export of our gas, the industry argued that it was unfair to change the rules after they had made their investments. This has resulted in a prolonged political fight, delayed reform, and billions in revenue forgone. Australians are not getting a fair share for their gas resources. **We cannot afford to repeat that mistake with AI.**

This is why we must make it clear now that a mechanism to ensure Australians share in the value of AI is coming, before AI companies have locked in their Australian operations under an assumption of zero obligation to the Australian public. We welcome AI investment in Australia. We want AI companies to build here and employ here. But the AI companies that profit in Australia – from Australian consumers, businesses, data, resources, and land – owe us a return. We must put the AI companies on notice that we will be developing rules to ensure that happens.

The specific mechanism should be developed through consultation with industry, economists, and the public, and could involve a tax on international technology and AI companies or a licensing regime where AI companies pay for the right to train or operate in Australia. Regardless, the time to act on this is now – we must not repeat the mistakes of the past.

Conclusion

The uncertainty around AI is real. No one knows exactly how powerful it will become, how quickly it will develop, or what its ultimate consequences will be for Australian workers, businesses and society. That uncertainty is not a reason to wait.

The government must **take an active approach to AI policy**. Otherwise, our future will be determined by the capabilities and risks of AI models 'grown' overseas.

The policies in this paper are a starting point. They are specific, practical steps that Australia should take now, regardless of how AI develops. **They are designed to be 'no regrets' policies**. Setting up the institutions that can keep pace with AI as it changes. Addressing the harms that are already happening to Australians. Preparing for the risks that are already foreseeable. Actively pursuing the opportunities that are already within reach. Ensuring that the benefits of AI reach all Australians, not just the handful of overseas technology companies positioned to capture them.

AI will continue to change rapidly. New models will emerge, bringing new risks and new opportunities that are not yet visible. The policies proposed here will not be the last word and they are not intended to be. Priorities will change and new policies will need consideration. But these policies are the right starting point – **and the time to act is now**.

Consultation

This is a discussion paper, not a final policy platform. It has been developed in consultation with Australian and international experts and input from my community in Curtin. Because of the scale and scope of the change ahead, it is vital that the Federal Government's approach to regulating AI is informed by as many perspectives as possible.

I want to hear from you.

Are there risks or opportunities this paper has missed? Are there policies that should be pursued but are not included here? Are there important details about design, implementation or trade-offs that deserve more attention? Do you disagree with the directions proposed? Whatever your ideas – be they big or small – I want to hear from you. We have an exciting and uncertain time ahead and we must seize the reins.

Scan the QR code or visit www.katechaney.com.au/ai_policy_paper to have your say.

Kate Chaney MP

FEDERAL MEMBER FOR CURTIN

CONTACT

E. kate.chaney.mp@aph.gov.au

W. katechaney.com.au

@katechaneyMP
@katechaneyMP
@ChaneyforCurtin
@katechaneymp.bsky.social
@katechaneymp

ELECTORATE OFFICE

Cnr The Boulevard and Floreat Ave
Floreat WA 6014
PO Box 186 Floreat WA 6014
T. 08 9388 0288

PARLIAMENTARY OFFICE

PO Box 6022
House of Representatives
Parliament House
Canberra ACT 2600
T. 02 6277 4655

AUTHORISED BY KATE CHANEY MP

Cnr The Boulevard and Floreat Ave
Floreat WA 6014

Printed on recycled paper and 100% carbon neutral.