

# Supporting Safe and Responsible AI

Tech Council of Australia Submission

August 2023



## Executive Summary

Artificial Intelligence (AI) is one of the most transformative technologies of our time, with the potential to offer our nation significant economic, social, environmental and strategic advantages. AI is already driving major advancements in areas such as healthcare, industry and public safety, and the future holds even greater potential. Generative AI alone could add between \$45 billion - \$115 billion a year to the Australian economy by 2030.

In a period of high inflation, and low productivity growth, an uplift from productivity enhancing technologies such as AI is vital to preserve Australia's standard of living and ensure that we maintain a competitive edge in the global landscape. Australia must not risk being left behind on AI.

As part of Australia's drive to grasp the significant opportunities of AI, the Tech Council fully supports the Government's goal to ensure AI development and deployment is safe and responsible, and we support the need for regulatory safeguards. While many common applications of AI are low-risk, there are risks and harms that can emerge that need to be managed.

In considering how to manage these risks, we need to understand that AI is the 21st century equivalent of manufacturing in the 20<sup>th</sup> century, in that it will unleash a widespread wave of productivity and economic enhancements. It will also involve a diverse set of practices and technologies that will make a broad array of products, that are used in a broad array of contexts, across almost every business and consumer environment.

As with manufactured products, it is important to recognise that AI products will vary considerably across different use cases. A clear, risk-based, and proportionate regulatory and governance approach will help us to capture the potential while effectively managing the different types of risks.

Australia's approach to AI governance and regulation needs to recognise that we already have a sound legal framework and regulatory model that covers the development and use of products. This includes technology-neutral laws (such as privacy, competition and consumer protection, anti-discrimination) and regulators that monitor compliance with them, as well as specialised laws and regulators in charge of particular industries or higher-risk areas (such as food safety or therapeutic and medical goods). It also includes an Australian and global model for developing standards that cover technology neutral and product / technology specific products, processes and behaviours.

There are no exemptions in these laws for AI development or deployment, and regulators are already starting to issue guidance for how they apply to AI development and deployment. Approaches such as standards development and assurance can also be applied to the safe and responsible development of AI products, as they are today in areas as important and diverse as vehicles, machinery and equipment, children's toys and medical products.

It is a strength of Australia that we start with such a stable, proven regulatory framework and trusted, well-known and expert regulators.

However, it is vital that the Australian Government provides clarity on how this model will apply to AI development. This includes clarifying the application of laws and standards to AI, coordinating efforts across these different areas to ensure Australia has a modern, trusted, and consistent regulatory framework and model for AI development and deployment, and identifying where the model needs reform or adaptation as new

technologies, products and uses emerge.

This can be done by setting out a regulatory strategy for AI in Australia, ensuring coherence with international frameworks, better coordinating efforts across government agencies and regulators, translating broad-based ethics into actionable guidance for organisations, establishing an expert body to provide advice on emerging technologies and issues, and undertaking targeted reform to address any genuine gaps.

The TCA does not support following the path of the EU in developing a standalone AI Act, particularly given the diversity and widespread application of AI technologies. Australia has never introduced a single Act to regulate “manufacturing”, but rather regulated the practices and products used in or produced by manufacturing and their associated risks through a system of technology-neutral and industry specific laws and standards, governed by expert regulators.

Diverting from this model, which has worked well for decades, would risk producing laws and rules that are quickly outdated, that hinder regulators' ability to respond to dynamic or domain specific developments (such as AI in medical devices versus the use of AI in a public safety or credit-checking context), and that are confusing for businesses and consumers accustomed to the current regulatory model and regulators. Extra layers of technology-focused regulation on top of our existing technology-neutral laws could drive capability and investment offshore, particularly if combined with even more restrictive laws in key policy areas like copyright/IP.

There is a more effective alternative model. Our submission outlines a proportionate 5-pillar Plan for Safe and Responsible AI in Australia, underpinned by 13 detailed recommendations.

## 1 Develop a regulatory strategy for AI

Australia needs a coherent and clear regulatory strategy for AI to help drive a consistent, practical and well-coordinated approach to AI regulation across the various strands of government. The strategy should be informed by expert advice and underpinned by best-practice regulatory principles, including the adoption of a risk-based approach and a drive towards international coherence and alignment.

### ***Recommendation 1: Prioritise the development of a new regulatory strategy for AI***

We back a proposal from the UTS Human Technology Institute to establish a strategy that sets out Australia's aims and objectives for AI regulation, promotes national coherence and efficiency, generally adopts a tech-neutral approach, and establishes consultative mechanisms with industry, technical experts and civil society. The strategy should help set expectations for our existing regulators and signal key best-practice regulatory approaches that Australia will adopt (as outlined in recommendations 2 and 3 below).

### ***Recommendation 2: Adopt a risk-based approach***

The new regulatory strategy should clearly signal the adoption of a risk-based approach that encourages regulators to focus efforts on higher-risk use-cases and applications of AI. It could include a framework to support a consistent approach across regulators to risk classification that is adaptable and flexible to evolutions in the technology. It should avoid wide-scale blanket bans or moratoria, which could prevent responsible and beneficial applications of AI to be further developed, instead clear guardrails for development of 'high-risk' use-cases should be

adopted (e.g. even technologies such as facial recognition can have responsible and beneficial applications, such as biometric identity verification to improve privacy and security).

***Recommendation 3: Ensure international coherence and interoperability***

As a small market with a globally-facing tech sector, Australia needs an approach to AI that supports coherence and interoperability with the global market, and we need to be deeply embedded in international standards setting processes. The regulatory strategy should:

- Ensure Australia's definitions and approach on AI policy continues to be informed by and interoperable with international frameworks (such as the work underway through ISO/IEC, IEEE, NIST and policy developments in key jurisdictions like the US, UK and Singapore).
- Lift Australia's engagement with international fora and standards-setting bodies (such as the ISO, IEEE, OECD, Global Partnership on AI, G7 Hiroshima Process and the UK's Global AI Safety Summit)
- Position Australia to take a regional leadership role in AI governance across the Asia-Pacific.
- Ensure that Trade Agreements support the development and trade in AI and digital products between Australia and our major trading partners.

## **2 Stand up an expert AI coordination model to uplift regulator capability**

While regulators have significant domain expertise in the areas they are responsible for (privacy, competition, discrimination etc.), they do not necessarily have technical in-house expertise in understanding emerging technologies like AI. An expert AI advisory and coordination model could support regulators to take an informed, coordinated, and consistent approach to regulating AI that leverages existing regulatory expertise, while encouraging the necessary uplift and capability building required for effective AI regulation.

***Recommendation 4: Establish an expert 'hub and spoke' AI coordination model***

An expert 'hub and spoke' coordination model would involve the creation of a new expert body (the hub), comprised of multidisciplinary experts on AI from a variety of relevant backgrounds, including technical, legal, academic and industry, from Australia and/or abroad. The body would work with individual regulators to build AI capability and provide expert technical advice to help interpret how existing laws apply to AI (underpinned by guiding cross-sectoral principles to ensure a consistent and coherent approach). This could be a temporary body, or it could be a form of ongoing commission, as some other stakeholders have recommended.

## **3 Deliver better regulatory guidance and enforcement of existing laws**

We do not need to re-write the whole rulebook for AI or other emerging technologies. Australia's model of core, technology-neutral laws, industry specific laws and standards and expert regulators (e.g. in therapeutic and medical goods) has worked well for decades.

Australia's technology-neutral laws already cover a wide range of areas relevant to AI risk, such as privacy, consumer protection, anti-discrimination, defamation, and intellectual property. Other jurisdictions such as the US and the UK have shown

how these laws can be interpreted and applied for AI, with the benefit of enhanced regulatory guidance and expertise.

***Recommendation 5: Clarify and enforce existing laws***

Support existing regulators to establish a considered, consultative and systematic process to provide specific guidance on how existing laws apply to the development and use of AI technologies.

***Recommendation 6: Do not establish a new AI Act or AI regulator***

We should not seek to replace or duplicate our existing regulatory frameworks and regulators with a new AI Act or regulator, which could lead to regulatory duplication, siloed expertise and capability loss across government.

## 4 Targeted review and reform to ensure our laws are fit for the digital age

We do not see a need to overhaul the whole regulatory framework to address challenges presented by AI. However, there are clear areas where further review or reform is needed to either update outdated legal frameworks or address genuinely unique issues or gaps related to AI.

***Recommendation 7: Coordinate with other review and reform processes across government, including privacy reform***

Ensure the Government's approach to AI regulation and governance is coordinated and integrated with other policy development processes currently underway across government. This includes proceeding with reform of Australia's privacy laws to modernise and improve consumer protections, while clarifying compliance requirements for industry and delivering greater coherence and interoperability with international laws (including in areas such as facial recognition technology, automated decision-making and consumer data use/consent).

***Recommendation 8: Review novel or grey areas of law***

Undertake separate, methodical review processes to examine novel or grey areas of law or standards presented by AI, including: (i) Intellectual property / copyright and generative AI, (ii) responsibilities and liabilities at different levels of the technology stack for AI systems; and (iii) approaches to foundation/frontier models.

## 5 Build Australia's AI skills, workforce, literacy and industry capabilities

While the debate about responsible AI often focuses on regulation and governance, we need to have an equal focus on other policy levers that will be critical to enabling trusted and responsible AI innovation.

Australia has built a thriving tech sector in the last two decades, with tech activity contributing \$167bn in GDP across the economy, and 935,000 tech jobs. Australia's area of strongest advantage in the tech sector is software products, in particular areas such as enterprise software, fintech products, quantum technologies and biotech. These are all areas that are highly complementary to AI, giving Australia a source of advantage in the global AI race. Further, because of the broad-based nature of AI technologies, they are and will be used across all industries in the economy, lifting growth and productivity.

To take advantage of this economic opportunity, Australia will need sustained investment in education, skilling/upskilling, industrial capability, research and digital literacy.

Australia also needs to plan strategically to enable research and to develop critical, national AI assets in both the public and private sectors, including in areas such as training datasets. Without this research and assets, Australia will fall behind in developing AI technologies and services, and may lose agency over future AI products.

Investment to develop our AI ecosystem, as well as the modernisation of government services, is also vital to ensure Australia remains competitive and confident in the development and deployment of AI.

***Recommendation 9: Continue building our tech talent pipeline and upskilling our workforce***

We need to continue with the reforms needed to help us address skills shortages and reach 1.2 million tech jobs by 2030 to ensure that Australia has access to the right expertise to support responsible AI development and deployment. This includes domestic education, skills and training reform (e.g. introduce digital apprenticeships to skill and upskill our workforce) and migration reform.

***Recommendation 10: Increase investment in AI research***

Support investment in research to advance Australia's AI capabilities, assets and technologies, and to create methods, tools and evaluation frameworks for responsible AI development.

***Recommendation 11: Increase investment to develop our AI start-up and scale-up ecosystem***

Leverage investment models like the National Reconstruction Fund to overcome funding gaps that are inhibiting the growth of AI start-ups and scale-ups here in Australia, remove barriers to specialist foreign investment, and ensure that our policy settings promote fair competition and open markets to ensure Australia remains an attractive place for global tech investment.

***Recommendation 12: Government as an exemplar of responsible AI***

Position the Australian Government as an exemplar on AI adoption and governance through the new Data and Digital Government Strategy, including by identifying and driving beneficial use cases across government, establishing best practice governance models and adopting best-practice standards.

***Recommendation 13: Increase digital literacy and responsible AI awareness***

Increase digital literacy and responsible AI awareness for citizens, SMEs, NGOs and other groups, including by supporting the development of guidelines, tools and assurance frameworks (including at industry-specific levels) that organisations will need to understand how to operationalise responsible AI.

## 1. Introduction

Artificial Intelligence (AI) is one of the most transformative technologies of our time. After decades of academic research, a broad array of AI applications already exist and are being used in a wide range of situations today. These applications will continue to become more sophisticated and ubiquitous as the field advances and AI adoption progresses.

While the deployment of consumer-facing AI technologies, particularly Generative AI, has sparked recent worldwide interest, AI technologies have been used in a diverse range of sectors for a considerable period of time. In finance, AI algorithms work to analyse market trends, optimise trading strategies, and detect fraudulent activities. In manufacturing, AI systems have helped automate repetitive tasks, optimise production processes, and ensure quality control. For critical sectors such as research, healthcare, education and public safety, AI is already driving significant advancements. It has revolutionised medical diagnostics and treatment to enable faster and more accurate detection of diseases and early intervention for patients. Just last year, AI models accelerated scientific progress to help generate new antibodies and aid hydrogen fusion. The ability to harness the transformative potential of AI is also growing in importance to help solve some of our most pressing societal and global issues, such as climate change, energy security and disaster response.

As AI models have become more affordable and higher performing, the economic potential has also increased. Generative AI alone has the potential to add between \$45b billion - \$115 billion a year to the Australian economy by 2030, if we create the right environment to enable AI creation and adoption.<sup>1</sup> Of this increased economic value, 70% will come from enhanced productivity in industries across the economy (e.g. by partially automating repetitive tasks that will free-up workers to focus on the more complex, creative and higher-value parts of their jobs), 20% from improved quality of outputs (e.g. by acting as a “co-pilot” alongside workers), and 10% in new products and services that will create jobs and businesses that were not previously possible.<sup>2</sup>

In a period of high inflation, and low productivity growth, an uplift from productivity enhancing technologies and innovation will be vital to preserve Australia's standard of living. It will also become increasingly vital to our offensive and defensive national security capabilities. The policy decisions we make today will have a major bearing on our capacity to realise this potential.

Countries such as the United States, China, UK and some European nations have already made significant strides in AI development, gaining an edge in the global market. Private investment in AI has also accelerated worldwide with global spend on AI estimated to be \$3 trillion by 2030.<sup>3</sup> While Australia has some leading AI capabilities in our start-ups, scale-ups and research sectors, our research shows that we significantly under-index on investing into our AI sector, attracting just 0.3% of total global VC investment.<sup>4</sup> The Productivity Commission has also shown that Australian businesses across the broader economy lag the global frontier in the uptake of advanced technologies like AI and data analytics.<sup>5</sup>

---

<sup>1</sup> Tech Council of Australia and Microsoft (2023), 'Australia's Generative AI opportunity'.

<sup>2</sup> Ibid.

<sup>3</sup> Statista, Artificial Intelligence (AI) market size worldwide in 2021 with a forecast until 2030.

<sup>4</sup> Pitchbook, Australia VC funding as share of global VC funding, 2017-2021.

<sup>5</sup> Productivity Commission (2023), Advancing Prosperity – 5-year Productivity Inquiry Report

Critically, a responsible, fair and proportionate regulatory framework is essential for Australia to both take advantage of the significant economic, social, environmental and strategic opportunities created by AI, and to manage potential risks associated with it.

As part of Australia's attempt to realise the enormous opportunities of AI, the Tech Council fully supports the Government's goal to ensure AI development and deployment is safe and responsible, and bounded by regulatory safeguards. While many common applications of AI are low-risk, there are real risks and harms that can emerge in certain use cases. These are categorised by the Human Technology Institute at UTS as follows:

- AI system failures (e.g. bias, discrimination, security failures);
- Malicious or misleading deployment (e.g. misleading systems, misinformation at scale, AI-powered cyber attacks); or,
- Overuse, inappropriate or reckless use (e.g. the erosion of privacy via inappropriate use of facial recognition technology, carbon costs due to excessive use).<sup>6</sup>

Importantly, there are unique use cases for different types of AI technologies, which will in turn pose different risks and challenges. For example, the risks of using AI technologies for internal business operations will be very different to use cases in public-facing government service delivery, a policing or justice context, or surgery. An AI system can also be developed for a beneficial purpose by one company, but deployed irresponsibly by another company with harmful results. We need guardrails and regulation, but the important question is how we do this.

In this time of high inflation and low productivity growth, Australia must not risk being left behind on AI. We need our policy settings to help foster a thriving AI ecosystem, attract and nurture top AI talent and investment, drive AI adoption and support the growth of an innovative AI sector.

Regulatory certainty surrounding the governance of AI through the provision of clear and predictable guidelines and governance frameworks is essential to enhance innovation, investment and adoption of AI in Australia, while also mitigating the harms and risks associated with AI development and deployment.

There are a range of policy levers Australia will need to utilise to move the arc of AI progress towards our economic and strategic goals. This includes standards setting, regulatory clarity, funding, skills, training and workforce development. Global competition to lead in AI is in full swing; the outcome of which will not only shape the economic landscape but will also influence societal and geopolitical dynamics in years to come.

---

<sup>6</sup> Solomon, L., & Davis, N., (2023) The State of AI Governance in Australia, Human Technology Institute, The University of Technology Sydney, p16.



## 2. Current regulatory landscape

There are five key points to understand about the existing landscape for safe and responsible AI in Australia:

- **AI systems in Australia are already subject to regulatory frameworks.** Australia has a robust regulatory model covering product development and deployment. This includes technology neutral laws and regulators, sector specific laws and regulators, and national and global standards processes.
- There are a range of existing technology-neutral laws that apply to the development and deployment of AI technologies, including privacy laws, anti-discrimination laws, competition and consumer laws, work health and safety laws, and IP/copyright laws. Directors' Duties under the Corporations Act are also relevant.
- **Sector specific guideline are also being prepared to support interpretation of the law for AI applications.** In addition to technology neutral laws, Australia already has a number of sector specific regulatory frameworks and regulators, particularly in areas of higher product risk. These regulators are already moving to clarify the rules for products under their jurisdiction that utilise AI. For example, the Therapeutic Goods Administration (TGA) has produced guidance on the regulation of software based medical devices, including AI. This helps software developers understand how the TGA interprets legislative requirements. This is a demonstration of how regulator guidance can help inform sector-specific governance.
- **There are strong existing and emerging international standards environment for AI technologies.** Australia is an active participant in national and international standards development. Standards are a critical tool in the regulation of the development and deployment of physical products, and can have similar value for AI based products. Already, there are important efforts underway for AI related standards. This includes the NIST AI Risk Management Framework, a range of relevant ISO/IEC standards (e.g. 22989, 23894, and 38507) and the forthcoming IEEE P2863, ISO/IEC 42001 and 42006 recommended practice for organisational governance of AI (the latter two, being developed by an Australian-led working group). Given that AI-specific regulatory and legislative frameworks are globally varied and inconsistent, standards can be highly effective measures to drive a coherent, safe and responsible international approach to emerging technologies while avoiding added costs and complexity for globally-facing businesses, like those in the tech sector.
- **There are a range of emerging responsible AI frameworks and practices developed and used by industry actors to manage risks related to AI systems.** These frameworks include Atlassian's Responsible Technology Principles, Adobe's Content Authenticity Initiative, and Google's Secure AI Framework (SAIF) and SEEK's Responsible AI Framework. There are also a number of technical governance mechanisms used across the AI product development lifecycle. This includes pre- and post-deployment risk assessments, external risk assessments and auditing, model documentation and/or transparency notes, data provenance notes, red-teaming, the adoption of common technical standards, and monitoring and shared reporting mechanisms on vulnerabilities, system capabilities, limitations and use. (see *Appendix B* for further descriptions). Australia's AI Ethics Framework includes 8 voluntary principles that organisations can apply.<sup>7</sup>

---

<sup>7</sup> Department of Industry, Science and Resources. Australia's AI Ethics Principles.

### 3. Governance Gaps

While there are many frameworks already in place to drive safe and responsible AI, there are important gaps where there is a case for action by government and/or industry:

- **There is a lack of clarity from regulators around how existing laws (as outlined above) apply to AI systems and how they will be enforced.** This regulatory uncertainty can hinder the positive adoption and development of AI in Australia, limiting our ability to capture the benefits of AI.
- **There are also some fundamental legal issues that need to be further clarified.** This includes where responsibility and accountability sits across the tech stack/product development lifecycle/supply chain for any given AI system, including identifying the appropriate roles (developers, deployers, data suppliers, end-users etc.); the corresponding governance responsibilities for high-risk use cases; as well as processes for operationalising explainability and transparency of AI systems.
- **Australia lacks a model to help coordinate AI regulation and policy across Government,** which creates a risk of disjointed, incoherent or even inconsistent requirements. There is also no formal source of expert advice to Government to inform its approach.
- **There are important areas of law where Australia has not modernised,** which have left our frameworks outdated compared to international norms – this includes privacy law and intellectual property / copyright law.
- **The public sector arguably lags behind the private sector** in driving mature and responsible AI governance practices internally.
- **Australia is also not as engaged in international standards setting processes for AI as it needs to be.**
- **There are major gaps in Australia's tech workforce, AI literacy, and in the domestic tech funding environment** which will inhibit our capacity to deliver responsible and trusted AI innovation in our country.

## 4. The Tech Council's 5 Pillar Plan for Safe and Responsible AI in Australia - Detailed Recommendations

### 1 Develop a regulatory strategy for AI

Australia needs a coherent and clear regulatory strategy for AI to help drive a consistent, practical and well-coordinated approach to AI regulation across the various strands of government. The strategy should be informed by expert advice and underpinned by best-practice regulatory principles, including the adoption of a risk-based approach and a drive towards international coherence and alignment.

#### ***Recommendation 1: Prioritise the development of a new regulatory strategy for AI***

We back a proposal from the UTS Human Technology Institute to establish a strategy that sets out Australia's aims and objectives for AI regulation, promotes national coherence and efficiency, generally adopts a tech-neutral approach, and establishes consultative mechanisms with industry, technical experts and civil society. The strategy should help set expectations for our existing regulators and signal key best-practice regulatory approaches that Australia will adopt (as outlined in recommendations 2 and 3 below).

#### ***Recommendation 2: Adopt a risk-based approach***

The new regulatory strategy should clearly signal the adoption of a risk-based approach that encourages regulators to focus efforts on higher-risk use-cases and applications of AI. It could include a framework to support a consistent approach across regulators to risk classification that is adaptable and flexible to evolutions in the technology. It should avoid wide-scale blanket bans or moratoria, which could prevent responsible and beneficial applications of AI to be further developed, instead clear guardrails for development of 'high-risk' use-cases should be adopted (e.g. even technologies such as facial recognition can have responsible and beneficial applications, such as biometric identity verification to improve privacy and security).

#### ***Recommendation 3: Ensure international coherence and interoperability***

As a small market with a globally-facing tech sector, Australia needs an approach to AI that supports coherence and interoperability with the global market, and we need to be deeply embedded in international standards setting processes. The regulatory strategy should:

- Ensure Australia's definitions and approach on AI policy continues to be informed by and interoperable with international frameworks (such as the work underway through ISO/IEC, IEEE, NIST and policy developments in key jurisdictions like the US, UK and Singapore).
- Lift Australia's engagement with international fora and standards-setting bodies (such as the ISO, IEEE, OECD, Global Partnership on AI, G7 Hiroshima Process and the UK's Global AI Safety Summit)
- Position Australia to take a regional leadership role in AI governance across the Asia-Pacific.
- Ensure that Trade Agreements support the development and trade in AI and digital products between Australia and our major trading partners.

## Recommendation 1: Prioritise the development of a new regulatory strategy for AI

- 1.1 One of the most important roles the Government can play in creating an effective regulatory framework for responsible AI in Australia is providing national leadership to drive coherence, capability and consistency across our various regulators and regulatory frameworks.
- 1.2 The Tech Council therefore supports a proposal from the UTS Human Technology Institute for the Government to establish an AI regulatory strategy for Australia..
- 1.3 A regulatory strategy should have a clear focus on international interoperability and alignment. These elements will help drive a best-practice approach to AI regulation across Government and set common expectations and principles around how our existing regulators should be approaching AI.

## Recommendation 2: Adopt a risk-based approach

- 2.1 AI is a cross-cutting technology and different use-cases and applications of AI across various industries and sectors will come with different harms, risks, and challenges. These will in turn, present varying degrees of regulatory challenge. A “one-size-fits-all” approach to AI regulation will over-fit in some instances and under-fit in others. Managing this complexity requires a nuanced approach that is sufficiently adaptable, flexible, targeted, and outcomes-based.
- 2.2 A risk-based approach is the appropriate regulatory approach to adopt for AI. It balances the need for governance and innovation, while crucially acknowledging the differences in context and use-cases in which AI applications are deployed based on their system outcomes and effects.
  - 2.2.1 The use of AI for low-level automation tasks to optimise efficiencies for business operations will have a significantly different impact and should be treated distinctly from an AI system deployed for use-cases in critical infrastructure, national security, or a policing and justice context, for instance.
- 2.3 This approach also aligns with international approaches to AI governance.<sup>8</sup> It includes the UK’s Pro-Innovation approach to AI regulation, Singapore’s Model AI Governance Framework for organisations adopting AI, the EU’s Artificial Intelligence Act, as well as broader frameworks including the World Economic Forum’s AI Governance Framework.
- 2.4 Risk-based frameworks have long been used by regulators and legislators to help define the risk detection, prevention, and mitigation steps that organisations should take in the context of hazards to society and the environment.<sup>9</sup> They share the general principle that risk management should:
  - i) target areas of where risks are greatest and,
  - ii) be proportionate and tailored to the degree and nature of risks.<sup>10</sup>

---

<sup>8</sup> Many other emerging international frameworks are risk-based given their subject of regulation, even if not explicitly stated. For example, draft laws in US states that these laws apply only to certain use cases are inherently risk-based because they have selected the use cases that they consider the highest risk and priority to regulate.

<sup>9</sup> Legislation pertaining to anti-money laundering and counter-terrorist financing, bribery and corruption, health and safety, food safety, anti-slavery, and -to certain extent- environmental due diligence legislation all place a strong emphasis on a risk-based approach.

<sup>10</sup> OECD (2022), ‘Translating a risk-based due diligence approach into law: Background note on Regulatory Developments concerning Due Diligence for Responsible Business Conduct.’

- 
- 2.5 A risk-based approach enables oversight measures to be tailored and ensures that governance is targeted, and that regulatory resources are allocated effectively and efficiently. Less restrictive oversight is assured for AI applications with lower risks, while robust governance mechanisms are reserved for those systems that need greater protection and oversight. This is important considering that AI systems are being used widely across society and the economy (and have been for a long time).
- 2.6 It is also important to note that a risk-based approach does not equate to zero risk. Instead of aiming for the absolute elimination of risk, this approach acknowledges that risk is an inherent part of any activity or process and in many cases, it is not feasible or practical to achieve zero risk.
- 2.7 Accordingly, there should be no distinction between private and public sector risk, as the core principles surrounding the classification of risk and risk-management are the same. Alignment on this approach also helps facilitate the exchange of best practices between sectors, while developing the maturity of risk-management frameworks for AI as a whole.
- 2.8 Risk-based approaches also reflect the ongoing, iterative nature of identifying and prioritising risks and impacts as they evolve and emerge, and require organisations to carry out regular risk assessments and ongoing monitoring.
- 2.9 This is especially helpful for AI models for a number of reasons. First, it may help in addressing issues in 'model drift', that is when an AI model's predictions or outputs start to become less accurate or reliable as it encounters new data after being deployed.<sup>11</sup> It is also useful for the development of early-stage models where the risk is unknown or unclear. As risk interpretations and tolerances are expected to change over time, these frequent changes can affect the effectiveness of regulatory tools. Risk management frameworks and standards are also a continual process and have greater adaptability and flexibility over instituting amendments to legislation, for example.
- 2.10 Small scale experiments in controlled environments should also be allowed without the need for prohibitive or large risk management processes to enable innovation, recognising that there is a low risk that the experiment will have significant impacts or consequences.
- 2.11 It is also important to note the potential risk of missed opportunities in not adopting AI. Failure to promote a conducive environment for AI innovation can result in the forgone benefits, the loss of competitiveness on the global stage, and hinder our nation's ability to harness the full potential of AI for social and economic growth.

#### Guidelines for high-risk AI use-cases and applications

- 2.12 In adopting a risk-based approach, the Government should ensure regulators focus their efforts on high-risk use cases and applications. This may include areas such as critical infrastructure, public-facing government service delivery, facial recognition, national defence, and security (this is not an exclusive list). These are some of the areas that could have the most significant risks for the safety and privacy of Australian citizens and should be prioritised.

---

<sup>11</sup> AI models are typically trained on data to learn patterns and relationships, however when an AI model is deployed and begins to interact with new and evolving data, the model's predictions or outcomes may change to become less reliable or accurate – this divergence is known as 'model drift'.

- i) Critical infrastructure – AI is increasingly integrated into energy, transportation, communication and healthcare systems. Disruption or compromise in these systems would have severe consequences on public safety, essential services and the economy. Safe, trusted and responsible deployment of AI systems in these areas is crucial. We note that the Government may consider some of the sectors covered in the *Security of Critical Infrastructure Act 2018* and that the Government should be cautious not to create new or duplicate provisions already covered by the Act.
- ii) Government service delivery – While Robodebt used data-matching and an automated-decision-making system (rather than AI), it is a demonstration of how important it is to ensure appropriate governance and oversight arrangements are in place for use of automated or AI based systems for public-facing government service delivery, particularly when used at scale. Ensuring the accuracy, transparency and ethical use of AI in these services is vital to maintain citizen trust in the government and the services provided.
- iii) Facial recognition – there are extensive applications of AI in law enforcement, border control, and identify verification. While FRT can be instrumental in enhancing security, it also raises a number of concerns related to privacy, civil liberties and misuse (see box overleaf)
- iv) National Defence and security – AI can play a pivotal role in intelligence gathering, cyber defence, and autonomous weapons systems. Adequate oversight of AI in these areas is essential to ensure responsible use, minimise the risks of unintended consequences and avoid potential harm to civilians, or the escalation of conflicts.

#### Do not initiate wide-scale blanket bans and moratoria

2.13 Initiating wide-scale blanket bans or moratoria without considering the context or use-case in which these technologies are applied could hinder the responsible and beneficial applications of AI. Such an approach overlooks valuable opportunities and benefits offered to consumers and society and clear guardrails surrounding high-risk use-cases and applications are needed.

2.14 As detailed above, a risk-based approach considers the technology as applied within its specific context, it does not assume zero risk, and levels of risk are met with their attendant regulatory, governance and oversight measures.

Facial recognition technology (FRT) is a prime example illustrating this dilemma. FRT raises concerns relating to privacy, misuse, unauthorised surveillance and the potential reinforcement of bias. If considering FRT being deployed for mass surveillance purposes, tracking or monitoring individuals without their knowledge or consent, use in commercial contexts for social profiling or data monetisation, it could be considered 'high-risk'.

However, there are also many beneficial use cases of FRT. This includes, for example, enabling convenient biometric identify verification by enabling customers the ability to unlock mobile phones, access personal and private accounts, conduct secure transactions quickly, reduce the reliance on passwords, and therefore enhancing both cybersecurity and user experience. As such, a thoughtful regulatory response that enables the risks to be considered with the benefits is crucial.<sup>12</sup>

<sup>12</sup> See Davis, N., Perry, L. & Santow, E. (2022) Facial Recognition Technology: Towards a model law, Human Technology Institute, The University of Technology Sydney.

### Recommendation 3: Ensure international coherence and interoperability

- 3.1 Australia is a small market with a global-facing tech sector. Given that Australia has a relatively small population of 26 million people, the domestic market for tech products and services is limited in scale and demand. As a result, Australian tech companies tend to be 'born global' and are often driven to grow and expand beyond borders to find larger customer bases and market opportunities.
- 3.2 Given the inherently global nature of technology as well as the potential reach of AI systems worldwide, international interoperability is critical for Australia's tech sector as well as many other parts of the economy.
- 3.3 Many Australian tech companies already 'benchmark' themselves by reference to global and overseas standards. In doing so Australian start-ups and scaleups reduce the friction in integrating with global technology markets. It also encourages foreign investment and investor confidence, as well as collaboration with technology companies abroad.
- 3.4 The degree to which Australia's regulatory response on AI is compatible and coherent with other jurisdictions has a significant impact on Australia's capacity to develop AI technologies as well as leverage and deploy them.
- 3.5 The Tech Council is encouraged to see that the discussion paper uses AI definitions that align with ISO standards. We support the alignment of definitions on 'artificial intelligence', 'machine learning', and 'algorithm' with international standard ISO/IEC 22989:2022. Using bespoke local definitions creates additional barriers to companies who want to engage in the global market.
- 3.6 We note the significant existing body of work already underway to develop international standards and governance frameworks for AI. This includes organisations including ISO/IEC, IEEE, the OECD, and the Global Partnership on AI (GPAI), as well as the current G7 Hiroshima AI process, the UK's upcoming Global AI Safety Summit and the ASEAN Guide on AI Governance and Ethics.
- 3.7 We strongly encourage that Australia takes a proactive role in actively embedding ourselves in these processes – to both contribute and benefit from expert discussions for the development of global technology regulations:
  - i) it is essential for Australia to stay informed on the latest technology developments and regulatory trends to leverage this knowledge and enhance our own national capability.
  - ii) it opens the door to collaboration with trusted cross-jurisdiction allies, through our Digital Trade Agreements for example; given that AI systems rely on large and diverse datasets, ensuring that cross-border data flows enable the continued development of emerging Australian AI ecosystem.
  - iii) finally, it ensures that Australian interests are considered and protected on the global stage.
- 3.8 We would also encourage Australia to take a regional leadership role in AI governance across the Asia-Pacific. There is an opportunity for Australia to lead a common AI regulatory framework across the Asia-Pacific that is aligned to other global initiatives. A first step could be through involvement in the development of the ASEAN Guide on AI Governance and Ethics.

- 3.9 While the growing influence of global standards does not mean Australia should adopt these provisions wholesale, it is crucial to emphasise that mechanisms for AI governance in Australia are interoperable and coherent for the benefit of Australian tech companies. The goal is to simplify the regulatory landscape for Australian tech companies to reduce the regulatory and compliance burdens, rather to increase their complexity.
- 3.10 We can also bolster participation by engaging industry representatives given their expertise in existing responsible AI practices, processes and frameworks that have already been developed (See current regulatory landscape and gaps, point 5).
- 3.11 There is also an opportunity to increase the awareness of Australia's involvement in the international standards processes in Government and the general public. There are lessons to be learnt from the role and importance of standards across sectors such as aviation, shipping, and pharmaceuticals (which could also be considered high-risk).

## 2 Stand up an expert AI coordination model to uplift regulator capability

While regulators have significant domain expertise in the areas they are responsible for (privacy, competition, discrimination etc.), they do not necessarily have technical in-house expertise in understanding emerging technologies like AI. An expert AI advisory and coordination model could support regulators to take an informed, coordinated, and consistent approach to regulating AI that leverages existing regulatory expertise, while encouraging the necessary uplift and capability building required for effective AI regulation.

### ***Recommendation 4: Establish an expert 'hub and spoke' AI coordination model***

An expert 'hub and spoke' coordination model would involve the creation of a new expert body (the hub), comprised of multidisciplinary experts on AI from a variety of relevant backgrounds, including technical, legal, academic and industry, from Australia and/or abroad. The body would work with individual regulators to build AI capability and provide expert technical advice to help interpret how existing laws apply to AI (underpinned by guiding cross-sectoral principles to ensure a consistent and coherent approach). This could be a temporary body, or it could be a form of ongoing commission, as some other stakeholders have recommended.

### Recommendation 4: Establish an expert 'hub and spoke' AI coordination model

- 4.1 Australia's existing suite of regulators have significant domain expertise that has been built up over a long period of time – whether that be the Office of the Australian Information Commissioner on privacy issues, the Australian Competition and Consumer Commission on issues of consumer protection and transparency, the eSafety Commissioner on online harms, or the Australian Human Rights Commission on discrimination and breaches of human rights.
- 4.2 The Tech Council believes it's critical to preserve this knowledge and capability, while enhancing the capacity of existing regulators to engage in digital technology issues, rather than creating separate standalone regulators for particular technologies like AI.
- 4.3 However, we also recognise there are challenges in uplifting the knowledge and capability of regulators on AI. This includes ensuring there is a relatively consistent and coherent approach across different regulators (to avoid sending confusing signals to the market which may deter investment and innovation).



- 4.4 We recommend that the Government stand-up an expert 'hub and spoke' coordination model to support existing regulators. This would include a core body comprised of technical experts in AI, as well as legal and policy experts from Australia and/or abroad. This group would also include industry experts to ensure that advice is workable and could be operationalised by businesses and organisations. By adopting a single point of coordinating, we can ensure that AI-related policies and decisions align with the broader vision and objectives set by Government.
- 4.5 Similar to the UK Government's central coordination function for AI or Singapore's Advisory Council on the ethical Use of AI and Data, this group would provide technical and expert advice to assist regulators in understanding AI systems and how existing laws can be applied to AI-related matters. The group would serve as the 'hub' which would be the focal point for aligning best practices to regulators and policy-makers who are the 'spokes'. It may be a temporary body, or it could be a more permanent fixture – the timeframes are not as important as the purpose of the body.
- 4.6 At the minimum, the group would be tasked with:
- i) building on our existing national AI ethics framework to develop a set of best practice principles for regulators on AI; and
  - ii) working with regulators to support the development of AI-focused regulatory guidance to interpret and apply existing bodies of law and legislation.
- 4.7 This model leverages the distinct expertise that existing general and sectoral regulators have developed. It would empower regulators to apply existing laws according to their regulatory domains which would enable more relevant, targeted, and context-specific regulatory application. It is possible that this group could complement the existing work of the Digital Platform Regulators Forum.

### 3 Deliver better regulatory guidance and enforcement of existing laws

We do not need to re-write the whole rulebook for AI or other emerging technologies. Australia's model of core, technology-neutral laws, industry specific laws and standards and expert regulators (e.g. in therapeutic and medical goods) has worked well for decades.

Australia's technology-neutral laws already cover a wide range of areas relevant to AI risk, such as privacy, consumer protection, anti-discrimination, defamation, and intellectual property. Other jurisdictions such as the US and the UK have shown how these laws can be interpreted and applied for AI, with the benefit of enhanced regulatory guidance and expertise.

***Recommendation 5: Clarify and enforce existing laws***

Support existing regulators to establish a considered, consultative and systematic process to provide specific guidance on how existing laws apply to the development and use of AI technologies.

***Recommendation 6: Do not establish a new AI Act or AI regulator***

We should not seek to replace or duplicate our existing regulatory frameworks and regulators with a new AI Act or regulator, which could lead to regulatory duplication, siloed expertise and capability loss across government.

## Recommendation 5: Clarify and enforce existing laws

- 5.1 The pervasiveness of AI across all domains necessitates that AI governance should evolve and build on existing legal frameworks, integrating AI within the existing regulatory landscape to ensure coherence and consistency.
- 5.2 The Tech Council strongly believes that our existing regulatory framework has the capacity and potential to apply to AI technologies, with the appropriate regulatory guidance that is informed by expert advice.
- 5.3 As the Discussion paper identifies, a combination of general and sector-specific laws, and standards, currently govern AI risks. Further work undertaken by the Human Technology Institute (HTI) at the University of Technology Sydney has also mapped AI harms to existing Australian laws. This includes the application of protection and privacy law, consumer law, competition law, corporations law, copyright law, online safety, discrimination law, administrative law, criminal law and the common law of tort and contract.
- 5.4 We strongly support the preservation of technology-neutral laws to ensure that our regulatory framework continues to adapt, evolve and be flexible. The adaptability and resilience of our legal system has been repeatedly tested and evolved to work in response to new business models and emerging technologies. In Australia, the cases of Trivago, Clearview AI, and Trkulja demonstrate how consumer law, privacy law, and defamation law have been effectively applied and interpreted to uphold crucial legal rights and protections in the digital age.
- 5.5 With this consideration, we do not currently perceive the need for comprehensive new laws for issues presented by AI, but instead, the focus should be on providing necessary guidance and fostering regulatory expertise and capability to apply existing laws.
- 5.6 As illustrative examples, we offer three areas where existing laws could be interpreted to work:<sup>13</sup>
- 5.6.1 **Consumer Law** – The Australian Consumer Law (ACL) could be triggered by the use of AI in ADM systems and chatbots; as well as where AI is the product or service (or part of) that is being marketed to the consumer. AI systems that engage in trade or commerce cannot engage in conduct that is misleading or deceptive, e.g. organisations must not misrepresent when an AI system is being used (that is, companies *providing* AI systems cannot engage in misleading and deceptive conduct when it comes to their AI systems under section 18 of the ACL).<sup>14</sup> and AI systems should not have 'safety defects' (given that AI systems themselves, are likely to constitute "goods" under the ACL, cannot contravene the consumer guarantees and other aspects of the ACL that apply to goods and services).
- 5.6.2 **Privacy Law** – The *Privacy Act 1988 (Cth)* stipulates the requirements for the collection, use and disclosure of personal information – and thus, the data being used in AI systems which applies to both data collected and used for training those systems and, where relevant, the inputs processed by and outputs generated by them, where such data contains personal information.; these obligations include

---

<sup>13</sup> Additional examples are provided in Solomon, L., & Davis, N., (2023) The State of AI Governance in Australia, Human Technology Institute, The University of Technology Sydney.

<sup>14</sup> The FTC's guidance on AI "claims" is drafted from the US perspective on their equivalent provisions to Australia's ACL section 18. That is, section 5 of the FTC Act on unfair and deceptive practices: <https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check>

taking 'reasonable steps' to implement organisational procedures, practices and systems to comply with the APPs for AI systems that 'learn' or develop over time, regular monitoring, and assessment.

- 5.6.3 **Anti-discrimination laws** – the risk of bias and discrimination in AI systems has been well documented, and so developers must actively monitor and assess their systems for bias or unfairness that would amount to direct or indirect discrimination.

### Recommendation 6: Do not establish a new AI Act or AI regulator

- 6.1 The Tech Council does not support the creation of a new overarching AI Act or single regulator for AI.
- 6.2 The current system of technology-neutral laws, industry specific regulation, and standards has worked well for decades for the development and deployment of products in Australia.
- 6.3 AI, like manufacturing, encompasses a diverse array of technologies, products and use cases. For the same reasons Australia does not have a single act governing "manufacturing", it would be counter productive to have a single act governing AI.
- 6.4 Layering further technology-specific AI regulation on top of our existing technology-neutral laws could add further regulatory complexity and confusion, and if combined with more restrictive regulation in key policy areas such as copyright, would risk driving capability and investment offshore.
- 6.5 Further, just as with manufacturing, there are instances where the domain context in which AI is used dictates its risks and the right regulator model. This is why specialist regulators such as the Therapeutic Goods Administration (TGA) to regulate medical products, food safety authorities to enforce compliance with food standards, and the ACCC regulates product safety. In these areas, regulators should continue to be able to develop domain specific guidance and regulation.
- 6.6 The creation of a single regulator for AI would likely lead to siloed expertise and capability loss across government and regulators. It would discourage the necessary broader uplift, capability building and development of regulatory expertise for digital regulation matters. And, in the long-term hinder our overall ability to adapt and evolve our existing regulatory architecture in an ever-changing technological landscape.

## 4 Targeted review and reform to ensure our laws are fit for the digital age

We do not see a need to overhaul the whole regulatory framework to address challenges presented by AI. However, there are clear areas where further review or reform is needed to either update outdated legal frameworks or address genuinely unique issues or gaps related to AI.

### ***Recommendation 7: Coordinate with other review and reform processes across government, including privacy reform***

Ensure the Government's approach to AI regulation and governance is coordinated and integrated with other policy development processes currently underway across government. This includes proceeding with reform of Australia's privacy laws to modernise and improve consumer protections, while clarifying compliance requirements for industry and delivering greater coherence and interoperability with international laws (including in areas such as facial recognition technology, automated decision-making and consumer data use/consent).

---

***Recommendation 8: Review novel or grey areas of law***

Undertake separate, methodical review processes to examine novel or grey areas of law or standards presented by AI, including: (i) Intellectual property / copyright and generative AI, (ii) responsibilities and liabilities at different levels of the technology stack for AI systems; and (iii) approaches to foundation/frontier models.

**Recommendation 7: Coordinate with existing review and reform processes across government, including privacy reform**

- 7.1 There are broader existing reform and review processes underway across government related to AI. This includes the current Privacy Act reforms, electronic surveillance reforms, digital identity reforms, and other existing review and reform processes. There is also the current House Standing Committee on Employment, Education and Training Inquiry into the use of generative AI in the Australian education system, discovery work underway by IP Australia to explore the potential impact of generative AI on intellectual property rights, as well as a New South Wales and South Australian Parliamentary inquiries regarding the use of AI across their respective States.
- 7.2 We encourage that the development of Australia's AI governance framework be coordinated and aligned with other policy development processes across government at both federal and state levels.
- 7.3 We particularly encourage the Government to proceed with reform of Australia's outdated private laws in a way that enhances consumer protections, while reducing the compliance burden for industry and promoting greater international interoperability. We direct the Department to our submission on the Privacy Act Review final report, which includes recommendations on a wide range of issues, including facial recognition technology, automated decision-making, and data use, protection and consent.

**Recommendation 8: Review novel or grey areas of law**

- 8.1 There are certain areas of law that raise novel legal questions and would benefit from more detailed review, consultation and possible reform. The Tech Council has identified the following three areas:
- 8.1.1 **Intellectual property and copyright** – Copyright is a critical intellectual property protection and an important part of our economic system. However, we note that Australia's current copyright regime is amongst the most stringent in the world (particularly compared to leading innovation nations like the US, Singapore, Israel and South Korea) and is outdated in many regards, as highlighted in key reports over the last decade from the Productivity Commission and the Australian Law Reform Commission. It is also an area that lacks clarity around its application to AI, particularly generative AI, and machine learning systems. This may limit our potential to design, develop, and train AI/ML models locally, while also potentially posing barriers to inward AI investment and adoption. It may also pose barriers to the adoption of Australian culture, language and values being embedded in generative AI systems – risking Australia's voice being left out of a cornerstone new technology of the 21st century. More detailed consideration and consultation is needed around how our copyright framework (and IP framework more broadly) should apply to AI/ML systems and whether reforms are required. In doing so, we need to distinguish between how content may be used to train an AI model versus the outputs of that model.
- 8.1.2 **Liability for AI systems** – Further work needs to be undertaken to clarify the roles and responsibilities for actors across the tech stack/product development

lifecycle/supply chain for any given AI system. This involves identifying the appropriate roles (developers, deployers, data suppliers, end-users etc.) and mapping AI governance actions accordingly. This work could be embedded in the remit of the proposed new export AI coordination model.

- 8.1.3 **Foundation/Frontier models** – These models are increasingly critical to the future AI ecosystem, serving as the backbone of many AI applications. Not only do these models underpin AI-powered products and services and also experiments in applying models to new use cases, they have the potential to push the boundaries of AI capabilities which involve novel approaches and experimental techniques that may have far-reaching societal impacts. Given their primacy, governance on these models would benefit from industry, researcher and policy collaboration in international standards fora. (see also 6.5, i-iii below).<sup>15</sup>

## 5 Build Australia's AI skills, workforce, literacy and industry capabilities

While the debate about responsible AI often focuses on regulation and governance, we need to have an equal focus on other policy levers that will be critical to enabling trusted and responsible AI innovation.

Australia has built a thriving tech sector in the last two decades, with tech activity contributing \$167bn in GDP across the economy, and 935,000 tech jobs. Australia's area of strongest advantage in the tech sector is software products, in particular areas such as enterprise software, fintech products, quantum technologies and biotech. These are all areas that are highly complementary to AI, giving Australia a source of advantage in the global AI race. Further, because of the broad-based nature of AI technologies, they are and will be used across all industries in the economy, lifting growth and productivity.

To take advantage of this economic opportunity, Australia will need sustained investment in education, skilling/upskilling, industrial capability, research and digital literacy.

Australia also needs to plan strategically to enable research and to develop critical, national AI assets in both the public and private sectors, including in areas such as training datasets. Without this research and assets, Australia will fall behind in developing AI technologies and services, and may lose agency over future AI products.

Investment to develop our AI ecosystem, as well as the modernisation of government services, is also vital to ensure Australia remains competitive and confident in the development and deployment of AI.

### ***Recommendation 9: Continue building our tech talent pipeline and upskilling our workforce***

We need to continue with the reforms needed to help us address skills shortages and reach 1.2 million tech jobs by 2030 to ensure that Australia has access to the right expertise to support responsible AI development and deployment. This includes domestic education, skills and training reform (e.g. introduce digital apprenticeships to skill and upskill our workforce) and migration reform.

---

<sup>15</sup> The formation of the Frontier Model Forum in late-July, an industry organisation founded by Microsoft, Anthropic, Google, and OpenAI, is one example of how industry is collaborating on sharing knowledge on the development of frontier models to ensure that AI advances responsibly and safely.

---

***Recommendation 10: Increase investment in AI research***

Support investment in research to advance Australia's AI capabilities, assets and technologies, and to create methods, tools and evaluation frameworks for responsible AI development.

---

***Recommendation 11: Increase investment to develop our AI start-up and scale-up ecosystem***

Leverage investment models like the National Reconstruction Fund to overcome funding gaps that are inhibiting the growth of AI start-ups and scale-ups here in Australia, remove barriers to specialist foreign investment, and ensure that our policy settings promote fair competition and open markets to ensure Australia remains an attractive place for global tech investment.

---

***Recommendation 12: Government as an exemplar of responsible AI***

Position the Australian Government as an exemplar on AI adoption and governance through the new Data and Digital Government Strategy, including by identifying and driving beneficial use cases across government, establishing best practice governance models and adopting best-practice standards.

---

***Recommendation 13: Increase digital literacy and responsible AI awareness***

Increase digital literacy and responsible AI awareness for citizens, SMEs, NGOs and other groups, including by supporting the development of guidelines, tools and assurance frameworks (including at industry-specific levels) that organisations will need to understand how to operationalise responsible AI.

## **Recommendation 9: Continue building our tech talent pipeline and upskilling our workforce**

- 9.1 Continuing to cultivate and expand Australia's tech workforce pipeline is crucial to address skills shortages and equip Australia with the necessary expertise to drive responsible AI development and deployment.
- 9.2 Our domestic education, skills and training system is a core part of growing our tech workforce. Curriculum adjustments could incorporate a greater number of foundational tech subjects at all levels - from primary school, to higher education institutions - to help foster a culture of digital literacy and develop the skills that our next generation needs to thrive in the 21<sup>st</sup> century. This includes ensuring we support students to engage with responsible use of generative AI technologies, rather than instituting bans or other wide-scale restrictions that will only inhibit students from developing skills in these technologies. Investment in scholarships, internships and digital apprenticeships can also incentivise young Australians to pursue technology and AI-related education programs to contribute to our skilled workforce.
- 9.3 Migration reform also plays a critical role in addressing our immediate skills shortages, in Australia's tech workforce and filling gaps in experienced workers, who play a key role in on-the-job training and mentoring local workers to develop Australia's future tech workforce. Skilled migrants bring the latest understanding of technology into organisations across Australia.
  - 9.3.1 This is true for the 100,000 tech workers in US tech firms who pass on knowledge and best practice to their co-workers, customers, governments and non-profits, in addition to the 4,000 workers who leave US tech firms each year and go into local tech roles both in the direct and non-direct sectors, the public service or found or scale start-ups in Australia. These firms play an outsized role in developing

experienced tech talent and can also catalyse our junior workforce. For example, when an early-mid career worker can work alongside an experienced worker, it lifts the productivity of a more junior Australian worker by 2.6% because of the extra coaching and training they receive.<sup>16</sup>

9.4 There are a number of policy actions the Government could take to develop efficient and simple pathways for highly skilled workers. This includes:

- i) prioritising employer-sponsor skilled migration, with fast pathways to permanency and increased labour mobility;
- ii) streamlining arrangements for visa holders earning more than the average full-time salary for sponsoring employers and removing occupational lists and labour market testing;
- iii) committing to a service standard of visa processing within 10 days;
- iv) ensuring that Australia remains an attractive place for global tech companies to invest, in an increasingly competitive global market; and
- v) considering an incentive framework to encourage workers with AI skills to migrate or return to Australia.<sup>17</sup>

### Recommendation 10: Increase investment in AI research

10.1 Investment in research should be guided by a long-term vision to support and develop our AI capabilities. Many of the AI-enabled products and services existing today have their origins in decades-old, federally funded basic research programs and given the strategic importance of AI, Government could drive significant investment in our domestic research sector to:

- i) fund initiatives that prioritise fundamental and foundational AI capabilities. This includes areas like perception, knowledge representation, learning, reasoning, as well as advancements in improved hardware that are more likely to result in scientific and technical breakthroughs, with the benefit for scale-up and adoption;
- ii) investing in the development of methods, metrics, and tools for responsible AI governance. This includes research on effective models for human-AI collaboration and the operationalisation of key concepts like verifiability, accountability, fairness, and bias mitigation, for example.
- iii) understanding the societal risks and potential harms associated with AI models. This would involve inclusive and interdisciplinary research on the impacts of AI, theoretical work on understanding AI techniques and their emergent properties, as well as the advancement knowledge on how to design AI models and systems that are reliable, dependable, accurate, and safe;

10.2 The Australian Government could also consider developing shared national datasets, assets, and environments for AI training and testing under government stewardship. This initiative could leverage the vast amounts of high-quality public data that currently exists to create an engine for research, collaboration, and responsible AI development. It would be important for this to be attended by clear guidelines on data

<sup>16</sup> Tech Council, Microsoft and LinkedIn, (2023), 'Harnessing the Hidden Value: How US tech workers boost the growth of Australia's tech ecosystem'.

<sup>17</sup> Brain drain in AI experts in Australia has arguably been stronger than other verticals in Australian technology, due to the relative lack of AI roles in Australia and the high salaries on offer in the US.

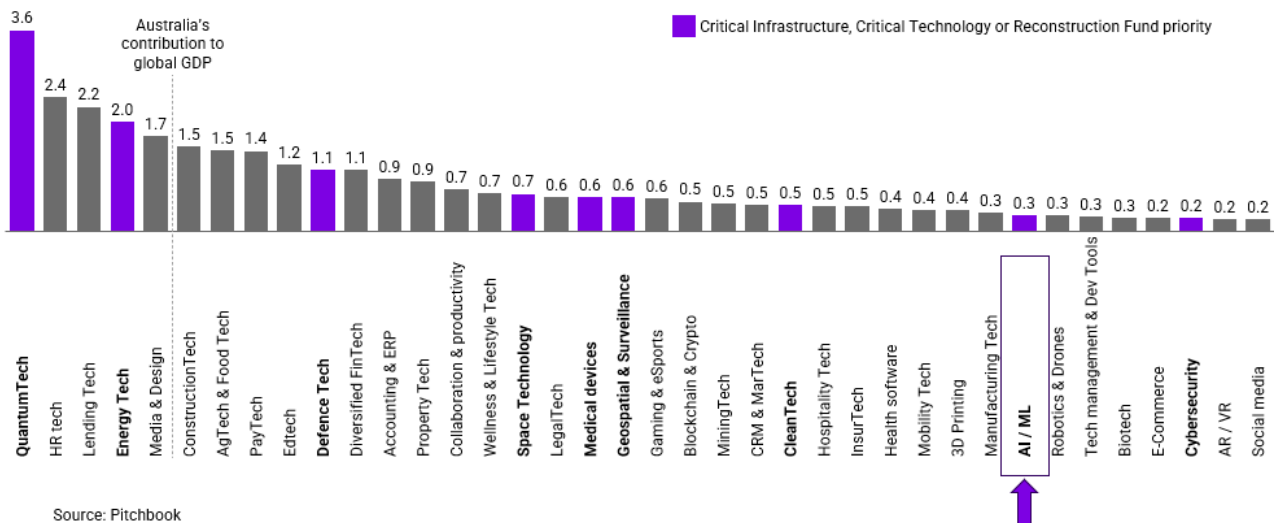
usage, privacy, and security, for example. Such an initiative could be housed within a government body, such as CSIRO/Data61.

## Recommendation 11: Increase investment to develop our AI start-up and scale-up ecosystem

11.1 Investment is vital to growing Australia's tech sector, which includes startups and scaleups developing or using AI-enabled products and services. It provides companies with the resources they need to develop, grow, and expand, as well as to attract top-tier talent in the market.

11.2 While Australia's venture capital market has matured considerably in recent years, our research shows that Australia generally under-indexes on investment into critical technology areas compared to global markets (despite often having a strong pipeline of start-ups in these areas) – see exhibit 1 below. Venture capital investment also remains well behind leading nations like the US, Singapore, UK and Canada on a per-capita basis, and we have a major funding gap in scale-up funding (from series B onwards).<sup>18</sup>

**Exhibit 1: Australian tech segment % share of global VC funding for that segment**



11.3 The Government could leverage existing investment models like the National Reconstruction Fund and the Industry Growth Program to address funding gaps and co-invest to support greater growth of our AI startups and scaleups.

11.4 The Government should also work to remove barriers to foreign investment to help overcome funding gaps that are inhibiting the growth of AI start-ups and scale-ups here in Australia and ensure that policy settings promote fair competition and open and transparent markets.

11.5 The presence of global companies in Australia provides additional benefits to the local ecosystem including skills transfer, productivity uplift, and fast-tracking career progression that help to ensure we have the necessary experience in our tech workforce – it is essential to ensure we retain open markets and promote fair competition.

<sup>18</sup> Tech Council of Australia (2023), 'Shots on Goal: A strategy for global success in tech'.



## Recommendation 12: Government as an exemplar of responsible AI

- 12.1 We encourage Government to uplift digital literacy and awareness across all departments and agencies and take a leading role as an exemplar of AI adoption and governance. This could be pursued as an objective of the new Data and Digital Government Strategy. Not only would this demonstrate the Government's commitment to innovation but would assist in improving public trust and confidence in AI systems.
- 12.2 Government could identify and drive beneficial use cases, informed by best practice approaches and governance models that are aligned to international standards, while ensuring cross-portfolio consistency throughout implementation. We encourage the Government to be bold in digitising and modernising its systems and operations. This will also require a shift away from legacy systems to modern, cloud-enabled systems.
- 12.3 The work currently undertaken by the NSW Government in responsible AI could also be a possible model to expand nationally. The NSW Government has created a state-based role for a Chief Data Scientist and has created an AI Assurance Framework which was developed with Australia's leading AI experts to apply across government agencies in NSW.

## Recommendation 13: Increase digital literacy and responsible AI awareness

- 13.1 Government plays a significant role in elevating digital literacy and fostering responsible AI awareness. Achieving this goal necessitates a multifaceted approach that includes education and awareness initiatives for the general public and accessible guidelines for organisations.
- 13.2 We encourage the Government to consider public-private partnerships to support digital literacy and responsible AI awareness. These arrangements bring together a range of stakeholders from government, industry, academic institutions and researchers, as well as civil society to develop initiatives that collectively encourage the responsible use of AI. A useful example of this is Singapore's Verify tool.
- 13.3 Education and awareness initiatives should focus on digital literacy as well as the safe and responsible use of technology more broadly. Public awareness campaigns and accessible citizen-friendly resources can provide useful knowledge on the importance of key topics such as data privacy, informed consent, cybersecurity, and others.
- 13.4 Organisations using AI-systems can equally benefit from education and awareness programs, as well as practical resources, toolkits, assurance guidelines and frameworks to navigate responsible AI use.

## Appendix A: TCA's guiding principles for regulatory design

The TCA accordingly recommends the following five guiding principles for best practice policy development in the digital economy:

- **Informed and coordinated** – technology regulation and policy development inherently addresses novel concepts and issues. For this to be effective, it requires us to have sufficient time, stakeholder input, and expertise to make informed policy decisions. Rigorous analysis and industry engagement, with thoughtful consideration of the interrelationships with other policies and regulation, helps us avoid the pitfalls of technical infeasibility and enhances regulatory compliance.
- **Proportionate** – a risk-based approach targeted at clearly defined problems enables regulation to achieve the objectives that are sought, while also avoiding unintended consequences such as increasing barriers to entry for others, or inadvertently capturing other parts of the tech sector.
- **Timely** – premature regulatory intervention can disproportionately impact emerging startups, business models, and technologies. To ensure Australia maintains a competitive place in the global market, we should be proactive in considering a range of potential policy levers, ensure that industry is given appropriate clarity and guidance, while enabling the appropriate opportunity and space for innovation.
- **Consistent and interoperable** – the technology industry is global by nature and few policy questions are unique to Australia. Regulation should consider and align, where appropriate, with domestic and global regulation to strive towards harmonisation and interoperability.
- **Has a bias to innovation and growth** – becoming a leading digital economy means that Australia should aim to encourage the responsible and early introduction and deployment of technology, this means avoiding prescriptive technical requirements that may become quickly outdated or inhibit innovation.

## Appendix B: Techniques and mechanisms for responsible AI

*^Note this is a non-exhaustive list of techniques used by industry*

Mechanism	Description
<b>Pre-deployment risk assessment</b>	<p>This occurs before the AI system is put into active use. It involves a comprehensive analysis of potential risks and challenges associated with the AI system's design, development, and planned usage. The primary goal is to proactively identify and address issues that may arise during deployment. This includes:</p> <ul style="list-style-type: none"> <li>- <b>Technical Risks:</b> Identifying vulnerabilities, biases, and limitations in the AI model, ensuring robustness, and addressing potential privacy concerns.</li> <li>- <b>Ethical Risks:</b> Evaluating the potential impact of the AI system on individuals, society, and vulnerable groups, and ensuring fairness, transparency, and accountability in decision-making.</li> <li>- <b>Legal and Regulatory Risks:</b> Ensuring compliance with relevant laws and regulations, such as data protection and anti-discrimination laws. <b>Operational Risks:</b></li> </ul>

	<p>Identifying potential disruptions, scalability challenges, and integration issues that may arise during deployment.</p>
<p><b>Post-deployment risk assessment</b></p>	<p>This takes place after the AI system has been deployed in a real-world environment. The purpose is to continuously monitor the system's performance, gather feedback, and address any new risks that emerge during operation. Key aspects of post-deployment risk assessment include:</p> <ul style="list-style-type: none"> <li>- <b>Monitoring and Feedback:</b> Continuously monitoring the AI system's behaviour, collecting user feedback, and identifying any unintended consequences or biases that may arise in real-world scenarios.</li> <li>- <b>Adaptation:</b> Making necessary adjustments to the AI system based on real-world data and feedback to mitigate risks and improve performance.</li> <li>- <b>Legal and Ethical Compliance:</b> Ensuring ongoing compliance with evolving laws, regulations, and ethical standards.</li> <li>- <b>Crisis Management:</b> Developing plans to handle unexpected issues, such as security breaches or major ethical concerns that may arise during operation.</li> </ul>
<p><b>External risk assessments and third-party auditing</b></p>	<p>External risk assessment and auditing for AI systems involve independent evaluations conducted by external third-parties, often experts or organisations not directly involved in the development or deployment of the AI system. This ensures an impartial evaluation of the AI system. These assessments aim to provide an objective and unbiased analysis of the AI system's risks, compliance with standards, ethical considerations, and overall performance.</p>
<p><b>Model documentation and/or transparency notes</b></p>	<p>This includes documentation or annotations to provide information about the model's design, development, usage and maintenance to understand the model's purpose, functionality, and operational considerations. Transparency notes are user-facing notes to provide insight into the workings of an AI system. Both may include information on aspects such as:</p> <ul style="list-style-type: none"> <li>- <b>Model architecture:</b> the components of the AI system including number and size of layers, types of layers (input, hidden, output), architecture designs or variants, etc;</li> <li>- <b>Training data:</b> descriptions of the data used to train the AI model, including data sources, size, quality, and note of any potential biases;</li> <li>- <b>Training process:</b> including optimisation techniques, loss functions, parameters and hyperparameters used;</li> </ul>

	<ul style="list-style-type: none"> <li>- Pre-processing and transformation: records of any data-pre-processing steps, such as data normalisation, augmentation, or feature engineering.</li> <li>- Model outputs: to detail how the AI system makes decisions or predictions, including the confidence or probability score, or decision thresholds applied;</li> <li>- Evaluation metrics: to evaluate the performance of an AI system, both during development and after deployment;</li> <li>- Capabilities and characteristics: including key functions and details on system behaviour;</li> <li>- System limitations and best practices: including known failure cases, scenarios where the model may not perform well and inversely, intended use-cases and considerations in choosing use-cases;</li> <li>- Updates and maintenance: information about how the system will be maintained, updated, and adapted to changing conditions; and,</li> <li>- Privacy and security: details on how the AI model handles privacy-sensitive data and security measures in place to protect against unauthorised access or data breaches.</li> </ul>
<p><b>Data provenance notes</b></p>	<p>This includes documentation that provides a record of the origins, history, and transformations applied to the data that is used to train, validate and test AI models. It includes data sources, the collection process, data pre-processing and transformation steps (as above), data quality, updates, version controls and any dependencies that are relied upon from external data sets, APIs, or third-party tools.</p>
<p><b>Red-teaming</b></p>	<p>Borrowed from cybersecurity, red-teaming for AI systems involves conducting simulated or adversarial testing on an AI system to identify vulnerabilities, weaknesses and potential areas for improvement. This can help assess the nature of unintended consequences in model behaviours, testing system's resilience to data poisoning, or other malicious activities.</p>
<p><b>Monitoring and shared reporting mechanisms</b></p>	<p>These mechanisms proactively support the mitigation of risks to ensure that the model operates safely and effectively. It includes the establishment of systems to monitor the AI's performance and behaviour in real- or near-time. This includes reporting on vulnerabilities, system capabilities, limitations and use It involves for example, collecting metrics, user feedback, and data on how the model is making decisions; monitoring for discrepancies across different demographic groups and taking corrective action to address bias; and monitoring data drift which may indicate the need for retraining and adjustment.</p>

---

	<p>The Frontier Model Forum founded by Anthropic, Google, Microsoft and Open AI is one example of an industry-led initiative that encourages reporting on AI models and facilitates information sharing on frontier model behaviour to support industry best practices and standards.</p>
<b>Adoption of common technical standards</b>	<p>As mentioned in our submission, there are a number of global efforts to develop standards on AI. Including the NIST AI Risk Management Framework, ISO/IEC standards (e.g. 22989, 23894, and 38507) and the forthcoming IEEE P2863, ISO 42001 and ISO 42006. These standards aim are a result of collaboration between various stakeholders including industry, governments, research institutions and technical AI practitioners. They cover different aspects including model development, security, transparency and auditing of AI systems and aim to promote uniformity and compatibility in best-practices.</p>