

**Submission
No 26**

INQUIRY INTO ARTIFICIAL INTELLIGENCE (AI) IN NEW SOUTH WALES

Organisation: Tech Council of Australia

Date Received: 20 October 2023

Inquiry into artificial intelligence (AI) in New South Wales

Tech Council of Australia Submission

October 2023



Executive Summary

The Tech Council welcomes the Parliamentary Committee's inquiry into Artificial Intelligence (AI) in New South Wales (NSW). We fully support the Committee's efforts to understand the use and impact of AI in NSW, its risks and challenges, and the economic and social opportunities.

The Tech Council is Australia's peak industry body for the tech sector. The Australian tech sector is a key pillar of the Australian economy, contributing \$167 billion to GDP per annum and employing over 935,000 people. This makes the tech sector equivalent to Australia's third largest industry, behind mining and banking, and Australia's seventh largest employing sector. NSW has the largest share of tech workers of any state or territory, with over 330,000 tech workers (around 35 per cent of the national total).

We represent over 160 companies from a diverse cross-section of Australia's tech sector, including companies working in business enterprise software, consumer software, telecommunications, fintech, venture capital and digital platform services. This includes companies that build and deploy AI-driven products and venture capital funds with investments in AI in NSW.

Our submission is presented in three parts:

- The first section provides an overview of the economic opportunities of AI,
- The second section deals with regulatory and governance approaches to manage the potential harms and risks of AI,
- Finally, the submission identifies a number of additional policy levers the NSW Government should consider to grasp the AI opportunity.

Key points from this submission:

- AI is one of the most transformative technologies of our time, offering significant economic, social, environmental, and strategic opportunities. It is also ubiquitous and has been used for many years to drive advancements in a wide range of industries and sectors, such as healthcare, accounting, energy and public safety.
- The potential for AI to address Australia's deep productivity challenge is immense. Generative AI alone is forecast to add between \$45 billion to \$115 billion a year to the whole Australian economy by 2030, largely through productivity enhancements, as well as the creation of new businesses and jobs.¹
- Realising the potential of AI will require a clear strategy, informed policy choices and targeted investment in skilling and upskilling the workforce, education, research, assets, adoption and the growth of new companies and industries.
- NSW has strong foundations to be a leader in AI. NSW's expansive tech workforce - the largest in the country - combined with a vibrant venture capital sector, innovative tech clusters and precincts (such as the Sydney Start-Up Hub, Tech Central, and Western Sydney Startup Hub), strong AI research capabilities, some of the country's most innovative tech companies and start-ups, and the Australian headquarters of a number of global tech firms, puts the state in a good position to be a part of the next wave of AI innovation. However, further policy change will be

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

needed to both capture the AI economic opportunity and address the risks from higher-risk use cases.

- A history of government support for the tech sector has led to NSW being one of the top 20 technology and innovation ecosystems in the world ranked by Startup Genome.² Supporting investment to continue developing NSW's tech and AI ecosystem is vital to ensure the state remains competitive not just in Australia, but in the broader Asia-Pacific region and globally.
- We support the goal of governments in Australia and across the world to ensure AI is safe and responsible. We also support the need for regulatory safeguards.
- Through this process, it is important to recognise that a number of laws and regulators already regulate AI development and deployment, including technology-neutral laws (covering a range of AI-risk areas such as privacy, consumer protection, anti-discrimination, defamation, and intellectual property), sector specific laws and regulators (e.g. medical products, financial services, food safety), and national and global standards processes.
- The most effective way to manage AI risks is to build on these existing laws and explain how they will apply to AI, while also drawing on international technical standards and industry best practices, not through introducing a new overarching AI Act or regulator.
- This is because AI will be used for a range of different purposes and products across almost every business and consumer environment - a one-size fits all approach won't work. We never tried to do this for "manufacturing" (which similarly applies to a wide range of products, use cases and consumer environments) and we shouldn't do it for AI.
- We should, however, undertake targeted review and reform in areas where there are genuine gaps or where there is a need for regulatory modernisation.
- We recommend the NSW Government ensure its governance approach to AI is nationally coordinated and that it supports existing regulators to clarify and enforce existing laws, and identify gaps. This could be supported by the establishment of an expert advisory and coordination model to support regulators, just like countries like the UK have done, and could also leverage NSW's AI Assurance Framework.

² StartupGenome (2023), Sydney, <https://startupgenome.com/ecosystems/sydney>

1. Overview of the AI economic opportunity

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. While the release of consumer-facing generative AI has sparked recent worldwide interest and debate, it is important to acknowledge that AI technologies have (and are) being used in a diverse range of sectors for a considerable period of time. These applications will continue to become more sophisticated and ubiquitous as the field advances and AI adoption progresses.

In public transport, AI is being used to provide smart ticking solutions, real time traffic flow updates, and optimise road safety. In manufacturing, AI applications help automate repetitive and labour intensive tasks, assist with work health and safety compliance, as well as detect anomalies in processes that could pose safety risks to consumers. In rural areas, farmers are using AI-precision tools to draw data from sensors and satellites to optimise planting, irrigation, and harvesting. For emergency and disaster resilience, early warning systems monitor and predict bushfire exposure areas to efficiently organise the deployment of emergency personnel and resources. In finance, algorithms work to analyse market trends, optimise trading strategies and detect fraudulent activities.

Just as connectivity to the internet and the adoption of technology has helped democratise access to knowledge, learning and education, especially through remote and rural learning, the ability to harness the transformative potential of AI is growing in importance to help solve some of our most pressing societal and global issues, such as climate change, energy security and disaster response.

For sectors such as research, healthcare, education and public safety, AI is already driving significant discoveries and advancements. Just last year, AI models accelerated scientific progress to aid the process of hydrogen fusion to transform energy. Astronomers are using AI to run simulations of dark matter and dark energy to better understand our universe. AI has helped solve one of the biggest problems in biological research by increasing the visibility over the structure of human proteins to accelerate future drug discovery.³ Based in Haymarket, Sydney-grown Harrison.ai, a clinician-founded medical and healthcare start-up founded in Sydney, is helping revolutionise medical diagnostics and treatment to enable faster and more accurate detection of diseases and early intervention for patients – with their product already in use by around one third of radiologists in Australia.

If we create the right environment to enable AI creation and adoption, generative AI alone has the potential to add between \$45b billion - \$115 billion a year to the whole Australian economy by 2030.⁴

To achieve these outcomes, we must make clear, informed, and conscious policy decisions. Achieving this will require a clear NSW Government strategy for AI, practical governance measures, and an uplift in investment in skills, assets, adoption and new companies and industries.

Workforce and skills implications

³ Scientists previously knew the 3D structures for just 17% proteins in the body, 98.5% of these structures (200 million proteins) are now able to be predictively mapped.

⁴ Tech Council of Australia and Microsoft (2023), 'Australia's Generative AI opportunity'.

The history of technological advancement has shown us that technology has a net positive impact on employment, while also playing an important role in driving productivity growth.

The biggest workforce challenge that AI presents is arguably how we skill and upskill our workforce to use and develop the technology productively and responsibly. As demand for AI-powered innovations grows, so will the demand for an AI-capable workforce. Of the increased economic value forecast to come from generative AI, 70% will come from enhanced productivity by partially automating repetitive tasks within a job 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 by using generative AI as a “co-pilot” to augment workers; and 10% in new products and services that will create jobs and businesses that were not previously possible.⁵

Australia already has major shortages in its tech workforce, particularly in technical and experienced roles such as software engineering and data scientists. We forecast that we will need another 600,000 people to enter the tech workforce to reach the Australian Government's and tech industry's shared goal of 1.2 million tech workers by 2030. We also need a general uplift in AI and digital literacy across the population.

State and territory governments will play a key role in fostering the right education and training settings to enable a workforce that can build, develop, and adopt AI. Achieving this goal will require a multifaceted approach that encourages AI use in the education system, supports reskilling and lifelong learning initiatives, establishes new training pathways such as digital apprenticeships, better recognises industry training initiatives, and involves collaboration with the federal Government to improve our migration system.

Supporting more Australians to get into the tech workforce will have positive economic and social benefits. Tech jobs are amongst the fastest-growing, best-paid and most flexible jobs in the economy. They have half the gender pay gap of other high-paying industries. According to SEEK data, tech sector jobs have the highest advertised salaries (\$132, 000) and the second highest advertised salary growth (6.0% year on year) of any industry.

International competition for AI leadership

Countries such as the United States, China, the United Kingdom, 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 2030.⁶

Our research, however, shows that Australia significantly under-indexes on investing into our AI sector, attracting just 0.3% of total global venture capital investment.⁷ The Productivity Commission has also recently reported that Australian businesses across the broader economy lag the global frontier in the uptake of advanced technologies like AI and data analytics.⁸

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. The policy decisions we make today will have a major bearing on our capacity to realise this potential. NSW needs to continue positioning itself to take advantage of this economic opportunity to ensure it is not left behind by other jurisdictions in Australia and across the world.

⁵ Ibid.

⁶ Statista, Artificial Intelligence (AI) market size worldwide in 2021 with a forecast until 2030.

⁷ Pitchbook, Australia VC funding as share of global VC funding, 2017-2021.

⁸ Productivity Commission (2023), Advancing Prosperity – 5-year Productivity Inquiry Report

New South Wales' strong foundations for AI leadership

NSW has an expansive tech workforce of over 330,000 which is the largest in the country, making up approximately 35 per cent of the national total.⁹ NSW's dynamic startup and scaleup ecosystem is supported by a robust venture capital sector that is instrumental in fuelling innovation and entrepreneurship, with most of Australia's largest venture capital firms based in NSW.

NSW has also launched a number of Australia's most innovative and globally successful tech companies including Atlassian, Canva, Afterpay, Airtasker, Employment Hero, WiseTech Global, and others. The NSW Government's history of ongoing support and investment in the tech sector has helped establish NSW one of the ecosystems and leading hubs for technological innovation in the Asia-pacific, with Sydney being ranked in the top 20 global ecosystems by the Startup Genome. This is exemplified by innovative tech clusters and precincts such as the Sydney Startup Hub, the Tech Central Scaleup Hub, and the Western Sydney Startup Hub. The vibrant ecosystem also helped to attract global multinational tech companies including Google, Microsoft, AWS and IBM who have established Australian headquarters in Sydney. In total, the tech ecosystem in NSW is valued at US\$78 billion, which is twice the global average.¹⁰

NSW is also home to a number of leading university and research institutes in AI. In addition to being home to CSIRO's National AI Centre, NSW has the Artificial Intelligence Institute (AII) based at the University of Technology Sydney (UTS), the UNSW AI Institute at UNSW, the Centre for Field Robotics which is one of the world's largest robotics institutes at the University of Sydney. NSW has also hatched a number of institutes focused on AI governance including the Allens Hub for Technology and Law at UNSW, the Gradient Institute at Sydney University, and Human Technology Institute at UTS.

New South Wales action on AI governance

NSW has also been an Australian leader in mechanisms for proactive governance of technology, including with the establishment of the Chief Data Scientist in 2015.

The NSW Government has achieved many 'firsts' including the first state in Australia to develop an AI strategy in September 2020 which was based on extensive consultation with the NSW tech and business community, academic leaders, ethics experts, and more than 1000 members of the public.¹¹ A year later, in 2021, the NSW AI Review Committee was established as the first of its kind in Australia. More recently, the NSW AI Assurance Framework came into effect in 2022 which also supports the set of mandatory ethical principles in for AI Ethics to guide NSW Government departments and agencies when developing, building, and implementing AI projects.

NSW's proactive stance on both investment, innovation, and governance is a core strength that has helped build the state's innovation ecosystem and could position it to be a leader in AI. We encourage the NSW Government to continue building on this approach.

⁹ Tech Council (2023), Tech Jobs Update, https://techcouncil.com.au/wp-content/uploads/2023/05/TechCouncil-Tech-Jobs-Update-May-2023_final-1.pdf

¹⁰ Startup Daily, <https://www.startupdaily.net/advice/opinion/nsw-budget-2023-startups-blueprint/>

¹¹ Digital NSW, Artificial Intelligence (AI), accessed at: <https://www.digital.nsw.gov.au/policy/artificial-intelligence>

2. AI Governance and regulatory frameworks

Understanding AI risks

Enabling regulatory certainty for AI 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.

While many common applications of AI are low-risk, there are potential risks that can emerge in certain circumstances. The Human Technology Institute at the University of Technology Sydney (UTS) has offered a useful categorisation of AI risks and harms which have been extracted here:

- 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).¹²

In considering how to manage and govern these risks, we need to understand that AI is the 21st century equivalent of manufacturing in the 20th 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. 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. 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.

To this end, there are five key points to understand about the existing regulatory landscape in NSW and Australia:

- **First**, it is a myth that AI is unregulated in Australia. AI systems in Australia are already subject to a number of regulatory and governance frameworks. Recognising the existence of our existing regulatory landscape, combined with established international technical standards, as well as industry best practice, can help achieve a holistic and operational approach to AI governance.
- **Second**, 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.
- **Third**, sector-specific regulatory frameworks and regulators also support the governance of AI applications. These exist particularly in areas of higher product

¹² Solomon, L., & Davis, N., (2023) The State of AI Governance in Australia, Human Technology Institute, The University of Technology Sydney, p16.

risk. Such 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.

- **Fourth**, there are strong existing and emerging international frameworks and standards for the development and governance of AI technologies. Standards are a critical tool in the regulation of the development and deployment of physical products and have similar value for AI based products. (See below for further discussion on standards).
- **Finally**, there are also range of industry best-practice AI governance mechanisms developed and used by leading industry actors to manage risks related to AI systems. They encompass techniques deployed across the AI product development lifecycle including pre- and post-deployment risk assessments, model documentation and/or transparency notes, data provenance notes, and shared reporting mechanisms. (See below for further details and **Appendix B** for a list of technical governance mechanisms and descriptions).

While there are many frameworks already in place to drive safe and responsible AI, there are some 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.
- While the NSW government has taken leading steps in coordinating AI development and deployment for government bodies and agencies, Australia lacks a model to help coordinate AI regulation and policy, which creates a risk of disjointed, incoherent or even inconsistent requirements.
- 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.
- More broadly, 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.

Recommendations to manage risk and harms relating to AI

Adopting a risk-based approach

A risk-based approach to AI governance is the appropriate approach to adopt to aid in mitigating the risks and harms relating to AI. A risk-based approach enables oversight measures to be tailored, 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). It is also important to note that a risk-based approach do not equate to zero risk.

Adopting a risk-based approach also 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. It also aligns with international approaches to AI governance.¹³

Risk-based approaches are iterative in identifying and prioritising risks and impacts as they evolve and emerge, requiring organisations to undertake regular risk assessments and monitor risks on a continuous basis. This is especially helpful for AI models. 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.¹⁴ 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. These frameworks and standards are also in continuous development and have greater adaptability and flexibility over initiating amendments to legislation, for example. For more on risk-based approaches, see *Appendix C*.

Leveraging international technical standards and industry-led approaches to AI governance

There are dozens of standards relating to AI from processes for governance and organisational oversight, to technical standards that assure the robustness of AI systems (as well as bias, transparency, explainability etc.). These include ISO/IEC (e.g. standards 22989, 23894, and 38507) and the forthcoming IEEE P2863. Standards can be highly effective measures to drive a coherent and interoperable international approach to regulating emerging technologies while supporting global-facing businesses, like those in the tech sector. Recently this year, Standards Australia established standard SA TR ISO/IEC 24027:2022 which adopted ISO/IEC TR 24027:2021 into the Australian standards landscape to address bias in relation to AI systems.

There are also regional and country-based standards developments, such as the US' NIST Framework and the ASEAN standards framework which will kick-off development later this year. ISO/IEC also has standards and guidance for the audit and certification of AI systems

¹³ This 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. 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.

¹⁴ 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'.

and IEEE hosts an ethics certification program. Other organisations such as the OECD, the Global Partnership on AI (GPAI), the current G7 Hiroshima AI process, the UK's upcoming Global AI Safety Summit, are all working to develop and align a global governance regime for AI. We recommend NSW to approach AI governance with these international technical standards in mind to leverage the thoughtful and expert work completed in these forums.

There are also a number of technical governance mechanisms used by leading industry actors that the NSW Government should be aware of. 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, as well as the work of the Frontier Model Forum.

These encompass mechanisms that are deployed across the AI product development lifecycle to assure the safe and responsible use of AI systems. Such measures include for example, 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 a list of technical governance mechanisms and descriptions.

The importance of interoperable regulatory and governance frameworks for NSW

Australia's tech sector is global facing, our relatively small population and domestic market for tech products and services means Australian tech companies must expand beyond our borders to find larger customer bases and market opportunities. Given the inherently global nature of technology, as well as the potential reach of AI systems worldwide, international interoperability for AI governance is critical for Australia's tech sector as well as many other parts of the economy.

Many Australian tech companies 'benchmark' themselves by reference to global standards, including those relating to AI. In doing so, our tech companies 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.

The degree to which any of Australia's regulatory responses to AI is compatible and coherent with the governance of other jurisdictions has a significant impact on Australia's capacity to develop AI technologies as well as leverage and deploy them.

A coordinated state-based assurance model for AI

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. 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.

If NSW is considering the adoption of new regulatory or governance regimes for AI, we would encourage it to leverage the expertise of its AI Review Committee and Chief Data Scientist, and adopt a technology-neutral, risk-based approach, one that is nationally coordinated with the work underway at the Federal level, and interoperable with international approaches.

- This includes coordinating with the Federal work underway on 'safe and responsible AI', as well as related processes such as the current Privacy Act reform process, electronic surveillance reforms, digital identity reforms, and the process being led by the Attorney-General on AI and IP rights.

For the same reasons Australia (and NSW) does not have a single act governing “manufacturing”, it would be counterproductive to have any single Act that governs AI. 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 out of the state.

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) regulate medical products, food safety authorities 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.

The creation of a 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.

Therefore, 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.

3. Further Policy Recommendations to realise NSW's AI Opportunity

There are a number of further policy actions that the NSW government can take to build the state's leadership on AI. These include further investment in AI research and development, continuing to develop NSW's AI start-up and scale-up ecosystem and enhancing digital literacy across the population.

Continue investing in AI research

The innovative applications of AI that we are seeing today have arrived only after decades of academic research. Many of the AI-enabled products and services existing today have their origins in decades-old, government funded basic research programs.

We encourage the NSW Government to pursue an ambitious plan for continued investment in AI research which should be guided by a long-term vision to support and develop AI capabilities. Given the strategic importance of AI, the NSW Government can continue to drive investment in:

- i) funding 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.

Investing to develop NSW's AI start-up and scale-up ecosystem

Government funding, especially non-dilutive funding and grants, helps early-stage businesses develop their first prototypes, gain their first customers, and encourages confidence for early investor rounds.

While we appreciate the NSW budget situation is tight, it will be vital to grow investment and jobs in new industries to get the budget back on track, and to develop technologies that help lift productivity. That is the only sustainable way to get inflation down and ease cost of living pressures. To do that, NSW needs to keep investing in the next generation of innovative companies and research in the state.

We welcome the new Innovation Blueprint process as it provides the opportunity to develop a sustainable future funding model. It is important the Government works with the industry as a priority to look at funding models in NSW, and how it can remain competitive versus other states, including by quickly clarifying what grant funding is available following recent budget changes. We are also committed to working with the State Government to make sure that support for office spaces for innovative businesses, such as at the Sydney Startup Hub, continues. The high cost of living and property prices impacts startups as well as households and small businesses. Ensuring affordable spaces to work is critical to keeping high potential new businesses in the state.

We encourage the NSW government to consider how to complement federal government initiatives, such as the National Reconstruction Fund and Industry Growth Program, and emulate successful models in other jurisdictions, e.g. LaunchVic.

Increase digital literacy and responsible AI awareness

As the NSW Government has already demonstrated through the establishment of the Sydney Startup Hub and Tech Central, public-private partnerships are an effective way to enhance knowledge transfer between industry, research and government sectors. This approach could also be channelled into providing broader digital literacy and AI awareness in organisations and the general public.

Education and awareness initiatives should focus on digital literacy as well as the safe and responsible use of technology more broadly. 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 which relate to AI.

Organisations using AI-systems can equally benefit from education and awareness programs, as well as practical resources, toolkits, assurance guidelines and frameworks that help organisations navigate the adoption and use of responsible AI.

The NSW Government as an exemplar of AI adoption

We finally encourage the NSW Government to continue uplifting digital literacy and awareness across all departments and agencies and take a leading role as an exemplar of AI adoption and governance – as it has done in recent years. Not only would this demonstrate NSW's commitment to innovation, but would assist in improving public trust and confidence in AI systems. This would involve identifying and driving 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 also encourage the NSW Government to be bold in digitising and modernising its systems and operations, including by replacing legacy systems.

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
Pre-deployment risk assessment	<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"> - Technical Risks: Identifying vulnerabilities, biases, and limitations in the AI model, ensuring robustness, and addressing potential privacy concerns. - Ethical Risks: Evaluating the potential impact of the AI system on individuals, society, and vulnerable groups, and ensuring fairness, transparency, and accountability in decision-making. - Legal and Regulatory Risks: Ensuring compliance with relevant laws and regulations, such as data protection and anti-discrimination laws. Operational Risks: Identifying potential disruptions, scalability challenges, and integration issues that may arise during deployment.
Post-deployment risk assessment	<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"> - Monitoring and Feedback: Continuously monitoring the AI system’s behaviour, collecting user feedback, and identifying any unintended consequences or biases that may arise in real-world scenarios. - Adaptation: Making necessary adjustments to the AI system based on real-world data and feedback to mitigate risks and improve performance. - Legal and Ethical Compliance: Ensuring ongoing compliance with evolving laws, regulations, and ethical standards. - Crisis Management: Developing plans to handle unexpected issues, such as security breaches or major ethical concerns that may arise during operation.
External risk assessments and third-party auditing	<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</p>

	<p>to provide an objective and unbiased analysis of the AI system's risks, compliance with standards, ethical considerations, and overall performance.</p>
<p>Model documentation and/or transparency notes</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"> - Model architecture: the components of the AI system including number and size of layers, types of layers (input, hidden, output), architecture designs or variants, etc; - Training data: descriptions of the data used to train the AI model, including data sources, size, quality, and note of any potential biases; - Training process: including optimisation techniques, loss functions, parameters and hyperparameters used; - Pre-processing and transformation: records of any data-pre-processing steps, such as data normalisation, augmentation, or feature engineering. - Model outputs: to detail how the AI system makes decisions or predictions, including the confidence or probability score, or decision thresholds applied; - Evaluation metrics: to evaluate the performance of an AI system, both during development and after deployment; - Capabilities and characteristics: including key functions and details on system behaviour; - 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; - Updates and maintenance: information about how the system will be maintained, updated, and adapted to changing conditions; and, - 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.
<p>Data provenance notes</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</p>

	<p>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>Red-teaming</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>Monitoring and shared reporting mechanisms</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>
<p>Adoption of common technical standards</p>	<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>

Appendix C: A primer on risk-based approaches for AI governance

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.¹⁵ 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.¹⁶

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).

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.

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.

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.

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.¹⁷ 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.

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.

¹⁵ 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.

¹⁶ OECD (2022), 'Translating a risk-based due diligence approach into law: Background note on Regulatory Developments concerning Due Diligence for Responsible Business Conduct.'

¹⁷ 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'.

In adopting a risk-based approach, the NSW Government should 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)

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.

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.