

We, Xaana Pty Ltd ("Xaana"), an Australian-owned and operated company of AI developers and architects, are based fully in Australia. We appreciate the opportunity to contribute to this significant inquiry into the adoption of artificial intelligence in our nation. As a homegrown organisation deeply invested in the Australian digital economy, we focus on developing AI technologies and solutions that cater specifically to Australian businesses and government agencies. Xaana delivers proven Turium Enterprise AI solutions in large Data Integration, Generative AI, and ERP Process Automation. Having achieved Authority to Operate in the Australian Federal Government, including for controlled unclassified information and classified protected networks. Xaana's Turium products are DISP, SOC2, and IRAP-Protected Certified. Supporting federal government agencies and large conglomerate companies such as Wesfarmers, Linfox, Bunnings, Kmart, and Shell Petroleum to push the boundaries of what AI can deliver where our innovation becomes their application. At our core, we embrace a strong ethos of Collaborative Intelligence (CINTEL), believing that technology should be used to maximise human capability, not replace it. Our commitment is unwavering: to develop safe Artificial General Intelligence (AGI) technology that serves all of humanity. By aligning our goals with the collective interests of our customers, partners, and the overall society, we consistently prioritise groundbreaking innovation over mere profit.

We have closely monitored and participated in discussions surrounding AI at various levels, including prior submissions to national inquiries. Our insights are informed by years of practical experience and continuous study and development of evolving AI technologies.

Global vs. Australia AI Adoption and Investment:

Global AI Investment:

- The global investment in AI is projected to approach \$200 billion by 2025, driven by substantial economic commitments, particularly in the U.S., where AI investment could peak at 2.5 to 4% of GDP¹.
- The allocation of \$116 million by the Australian federal government in the 2024 budget for AI technologies². While the Australian government's investment is a positive step towards fostering AI development, it is relatively modest compared to the larger amounts being directed towards AI globally. For example, in the United States, AI investment could approach \$100 billion alone by 2025³.
- Australia's AI adoption pace is notably slower; in 2022, only 24% of Australian businesses had deployed AI, compared to 39% in Singapore and 34% in Germany. This places Australia 13th out of 14 leading economies in terms of AI deployment⁴.

Impact on Workforce and Economy:

- Globally, AI is predicted to significantly enhance productivity and operational efficiencies. For example, generative AI investment surged to \$25.2 billion due to increased funding activities⁵.
- Australian AI adoption is notably slower compared to global standards. A report from BCG indicates that only about 30% of Australian organisations have succeeded in delivering digital

¹ Goldman Sachs, "AI investment forecast to approach \$200 billion globally by 2025," *Goldman Sachs*, Aug. 01, 2023.

<https://www.goldmansachs.com/intelligence/pages/ai-investment-forecast-to-approach-200-billion-globally-by-2025.html>

² Treasury, "July Economic and Fiscal Outlook," *budget.gov.au*, Jul. 23, 2020. <https://budget.gov.au/>

³ J. Sor, "AI could power the US economy as investment in the sector is poised to hit \$200 billion by 2025, a Goldman Sachs says," *Business Insider*, 2023. <https://www.businessinsider.com/ai-artificial-intelligence-stocks-investing-us-economy-outlook-goldman-sachs-2023-8>

⁴ ACS, "ACS Australia's digital pulse 2023," *www.acs.org.au*, 2023. <https://www.acs.org.au/insightsandpublications/reports-publications/digital-pulse-2023.html>

⁵ Stanford University, "The AI Index Report – Artificial Intelligence Index," *aiindex.stanford.edu*, 2023. <https://aiindex.stanford.edu/report/>



transformation, which is a critical step towards effective AI implementation⁶.

- Furthermore, Australia invests significantly less in AI compared to leaders like Israel, Singapore, and the US, and the country's AI maturity level lags behind the global average, with Australian companies self-scoring an average of 3.5 out of 10 compared to the global average of 4.3⁷.
- A critical barrier to AI adoption in Australia is the significant skill shortage in the tech sector. The Australian Computer Society (ACS) points out that nearly all Australian jobs will be transformed by AI whereby 11.2 million Australian workers will need to reskill, and \$16 billion cost to the Australian economy from tech skills shortage, demanding a concerted effort to reskill the workforce in line with future needs⁸.

Australia's Current Generative AI Adoption Picture

As a sovereign AI company deeply integrated within the Australian tech landscape, Xaana has observed firsthand the current dynamics of generative AI adoption in the country. The Australian Government has notably embarked on a significant trial of Microsoft 365 Copilot, engaging over 7,400 public servants across more than 50 agencies. This marks one of the first governments worldwide to deploy this generative AI service⁹.

Observations and Concerns from the Field:

- Collaborating predominantly with global tech giants like Microsoft has led to a monopolistic control over AI technologies. This dependence on foreign AI solutions stifles the growth and competitiveness of local AI firms and undermines the development of a sovereign AI industry in Australia.
- Trials involving significant data interactions with international firms raise serious concerns about data sovereignty. The use and storage of sensitive governmental data by entities governed by foreign laws could pose risks to national security and privacy.
- Heavy reliance on international AI technologies might impede Australia's economic independence. Developing local AI solutions would not only retain economic benefits within the country but also foster innovation tailored to Australia's unique requirements.
- The Australian government must implement robust policy and regulatory frameworks that encourage the growth of the domestic AI market, protect data sovereignty, and ensure that the benefits of AI technologies are widely distributed across the economy.

It has been noted that many foreign companies with Australian ABNs have million-dollar procurement contracts white labelled as intended for sovereign developments, which raises questions about the genuine 'sovereignty' of such initiatives.

Recently, there has been a move by Australian governments to develop their own generative AI solutions, such as the ChatGPT-like AI, rather than leveraging sovereign LLM solutions like Turium Zebra—our product.











⁶ P. Forth and S. Mohr, "The Success Formula for Australian Companies to Digitally Transform," *BCG Global*, Mar. 14, 2021. <https://www.bcg.com/publications/2021/digital-transformation-in-australia>

⁷ <https://www.bcg.com/about/people/experts/patrick-forth>, "Accelerating Australia's AI Adoption," *BCG Global*, May 02, 2023. <https://www.bcg.com/publications/2023/accelerating-australia-ai-adoption>

⁸ ACS, "ACS Australia's digital pulse 2023," www.acs.org.au, 2023. <https://www.acs.org.au/insightsandpublications/reports-publications/digital-pulse-2023.html>

⁹ "Australian Government collaboration with Microsoft on artificial intelligence | Prime Minister of Australia," www.pm.gov.au. <https://www.pm.gov.au/media/australian-government-collaboration-microsoft-artificial-intelligence>

Build In House vs Buy Generative AI

Criteria	Build In-House (Government Developed)	Buy Global LLM (International Vendors)	Buy Sovereign LLM 
 Skill Requirement	High demand for in-house AI expertise and continuous training.	Relies on international vendor expertise; less control over tech updates and skills development.	Leverages local industry expertise, enhancing local tech skills and development.
 Cost	High initial development and maintenance costs.	Potentially lower initial costs but may include higher operational costs and dependencies.	Competitive pricing structured for local markets
 Time to Deployment	Typically longer due to development cycles and testing phases.	Quicker to deploy as solutions are pre-developed; integration times vary.	Rapid deployment tailored to local systems and regulatory environments; faster integration.
 Customisation	Maximum customisation possible, tailored to specific government needs.	Limited customisation options; may not fully meet local requirements or preferences.	High customisation to meet specific local standards, regulations, and needs.
 Security and Compliance	Full control over compliance with national security and privacy standards.	Potential risks with compliance to local standards, depending on vendor's international focus.	Optimal compliance with national laws; easier legal and regulatory oversight.
 Innovation and Upgrades	Direct control over updates and innovation, but requires substantial ongoing investment.	Dependent on vendor's innovation cycle, which may not align with government's timeline or needs.	Collaborative innovation focused on national priorities and needs; more aligned updates.
 Economic Impact	Direct investment in local tech capabilities, but could be resource-intensive.	Primarily benefits the international vendor's economy.	Supports the national economy, potentially boosting local tech industry and job creation.
 Operational Risk	High due to potential for inexperience and technical challenges in emerging tech areas.	Dependency on external vendors for critical operations and support, which might affect response times and service quality.	Reduced risk through localised support and understanding of the domestic market and infrastructure.
 Data Sovereignty	Best control over data, ensuring all operations adhere strictly to national data laws.	Possible concerns over data being stored or processed under different legal jurisdictions.	Enhanced data control, with all operations and storage complying with national sovereignty laws.

Current Engagement Dynamics:

In our engagement with various Australian government agencies, the aftermath of the Robodebt Scheme has notably shaped our interactions and the broader AI adoption landscape.

Since we launched Turium Zebra ^{10 11} (Our Sovereign Large Language Model) in 2023, after years of research all the way back to 2019, there has been a noticeable increase in interest from Australian agencies wanting to understand and explore our AI solutions. We have been invited to numerous

¹⁰ Turium.AI, "Turium.AI Launching ZEBRA, The World's First Human AI Twin," [www.defenceconnect.com.au](https://www.defenceconnect.com.au/joint-capabilities/12948-turium-ai-launching-zebra-the-world-s-first-human-ai-twin), Oct. 12, 2023. <https://www.defenceconnect.com.au/joint-capabilities/12948-turium-ai-launching-zebra-the-world-s-first-human-ai-twin>

¹¹ A. McLaughlin, "Canberra-based company launches world's first human artificial intelligence twin," *Riotact*, 2023. <https://the-riotact.com/canberra-based-company-launches-worlds-first-human-artificial-intelligence-twin/717381> (accessed May 10, 2024).



meetings and demonstrations, indicating a genuine interest in the capabilities and potential improvements our technology can offer.

Despite the interest and positive feedback from demonstrations, there remains a significant hesitancy to commit to purchases. The shadow of the robodebt scandal looms large, with agencies being exceedingly cautious about how they implement new technologies that automate or significantly influence decision-making processes involving the public.

For those agencies that proceed with adopting our solutions, there's a specific condition: the term 'AI' is often requested to be downplayed or not mentioned at all in any public-facing or internal communications. Agencies are keen on leveraging the benefits of AI while minimising any association that might evoke past controversies.

How Xaana is playing a role in AI Adoption Education:

In response to the hesitancy surrounding AI adoption, Xaana Pty Ltd has initiated several strategic efforts aimed at educating and reassuring government agencies and the public about the benefits and safe use of AI technologies:

We have hosted a series of free webinars, bootcamps and forums specifically designed for federal government employees. These sessions aim to demystify AI technologies, discuss their potential benefits, and address common concerns related to AI implementation in government settings. The forums provide an open platform for dialogue, allowing government personnel to engage directly with AI experts and address their queries and concerns.

In collaboration with Torrens University, we are developing microcredentials focused on generative AI¹². These educational programs are designed to equip participants with the necessary knowledge and skills to understand and leverage generative AI in their respective roles. The curriculum includes practical applications of AI, ethical considerations, and guidelines for implementing AI solutions responsibly.

Recognising the need for comprehensive education on AI, we have established the Artificial General Institute (UAGI)¹³. This initiative focuses specifically on the public sector, offering courses and workshops that cover AI governance frameworks, ethical AI usage, and strategies for AI deployments.

Xaana Pty Ltd is proactively working to demystify AI and amplify its voice in the technology community through strategic partnerships and memberships. Here's how these relationships support our mission:

- By partnering with industry giants like NVIDIA, SAP, IBM, and Fujitsu, Xaana is harnessing their expertise to build and demonstrate practical AI applications and clarify common misconceptions about AI capabilities and limitations to their users. These partnerships help in presenting our real-world use cases that make AI more accessible and understandable to the public and the tech community.
- Australian Information Industry Association (AIIA): Membership in AIIA places Xaana in a position to influence industry standards and policies. It allows us to participate in public discussions, forums, and advocacy groups, helping to shape a more informed and balanced view of AI technology especially with the support of Simon Bush who has been a key influencer in our sovereign AI journey^{14 15}.

¹² Torrens University, "Torrens University and Xaana.AI shaping the future of AI education," www.torrens.edu.au, 2024.
<https://www.torrens.edu.au/stories/newsroom/university-news/torrens-university-and-xaana-ai-shaping-the-future-of-ai-education>

¹³ Xaana, "University of Artificial Intelligence UAGI | Xaana.AI," *Xaana*. <https://www.xaana.ai/uagi>

¹⁴ "MEDIA RELEASE," 2023. Accessed: May 10, 2024. [Online]. Available: <https://aiia.com.au/wp-content/uploads/2023/12/2023-12-08-AIIA-Media-Release-17M-Responsible-AI-Adopt-Program-to-support-Australian-Businesses.pdf>

¹⁵ "AIIA Member Directory," *AIIA*. <https://aiia.com.au/membership/we-keep-great-company/> (accessed May 10, 2024).



- Australian Computer Society (ACS): Through ACS, Xaana engages with the broader IT professional community, contributing to educational initiatives and policy discussions that aim to clarify the ethical and practical aspects of AI. This involvement also facilitates ongoing learning and development for our team, ensuring we remain at the forefront of AI advancements and discussions.
- Xaana is listed on the National Artificial Intelligence Centre, which is an initiative by CSIRO (Commonwealth Scientific and Industrial Research Organisation), Australia's national science agency in promoting and enabling AI adoption within Australia's AI ecosystem¹⁶.

Ethical Operational AI

According to us, the term "artificial intelligence" (AI) has become a buzzword, encompassing a vast and often ambiguous range of technologies. What was once described as "big data" or "predictive analytics" can now be readily rebranded as AI, blurring the lines between distinct concepts. Additionally, confusion arises from the tendency to conflate AI with automation.

This lack of clear definition creates a fertile ground for misleading claims. The promise of transformative AI advancements often falls short of reality, leading to disillusionment and concerns about ethical implications like algorithmic bias, accountability, and transparency. As a result, the field of AI finds itself in a critical state, struggling to live up to the inflated expectations.

Moving forward, a more critical and nuanced understanding of AI is essential. We must deconstruct the "AI" label and recognise the diverse range of technologies it encompasses. Only then can we move beyond the hype and focus on ethically and responsibly developing and deploying these technologies to address real-world challenges.

Has AI Ethics Lost Its Direction?

The discourse surrounding AI ethics has become an unfortunate echo chamber of lofty principles and abstract discussions, often offering little practical guidance for real-world application. This "ethics as theory" approach, prevalent in many AI ethics statements, resembles a check-the-box exercise, failing to address the complex ethical challenges faced by users and operators.

The sheer volume of these principles has fueled an industry dedicated to analysing them, further highlighting their inadequacy in addressing practical issues. This has sparked concerns about a legitimacy crisis surrounding AI ethics and begs critical questions:

- How can abstract principles translate into meaningful action and avoid being mere theoretical musings?
- While frameworks addressing algorithmic bias, accountability, and explainability are crucial, they represent a narrower focus. We argue against tunnel vision, urging the consideration of the entire system, not just isolated algorithms.

Our approach to technology ethics recognises that software platforms exist within a broader context, inextricably tied to their application, operational use, and surrounding data environment. This holistic perspective goes beyond the specific AI component and acknowledges the ethical implications of the entire system.

What is Xaana's Approach to Ethical AI?

If we start to take artificial intelligence more as tools for human use, we become better equipped to situate AI in appropriate framing contexts that recognise its critical features, constraints, liabilities, and dependencies:

¹⁶ CSIRO, "Xaana.ai," www.csiro.au. <https://www.csiro.au/en/research/technology-space/ai/ai-capabilities-directory/xaana>

We view AI models not as independent entities capable of miraculous solutions, but as tools within a larger system. Their capabilities are dependent on the supporting infrastructure, such as data quality, computational resources, and surrounding workflows. Additionally, AI models are vulnerable to errors if not carefully managed, maintained, and monitored.

Our notion of operational AI moves beyond the realm of theoretical ideas and academic exercises. We prioritise the integration of AI models into real-world scenarios, taking the full context of their deployment into account. This includes factors such as:

- Model inputs: Examining the quality and potential biases within the data used to train and operate the model.
- Users: Understanding the needs, skills, and limitations of those interacting with the AI system.
- Model outputs: Analysing the potential ramifications and unintended consequences of the AI's outputs.
- Consequences: Evaluating the real-world impact of the AI system on individuals, groups, and society as a whole.

We acknowledge the limitations of relying solely on fairness metrics and simplistic notions of data bias. Fairness is context-dependent, and a metric that may appear fair in one scenario might not translate effectively to another. Our approach focuses on building models that are fair and unbiased within the specific context of their intended use. This involves:

- Understanding the cultural, historical, and institutional background against which fairness needs to be evaluated.
- Recognising that all data exhibits some form of bias, the crucial question becomes which biases are acceptable or even necessary for the model to function correctly in its specific context.

We advocate for a holistic approach to data and model management, providing tools for continuous testing, evaluation, and improvement. This includes:

- Tracking the full provenance and lineage of data used in the model throughout its lifecycle.
- Structuring modelling efforts around a sensible relationship ontology that translates raw data elements into meaningful concepts based on the specific context.
- Implementing version control for changes made to data, models, parameters, and other elements of the system.
- Monitoring how dynamic environmental factors can affect usage and outcomes, ensuring ongoing model performance and reliability.
- Conducting continuous testing and evaluation, including data quality checks and integrity assessments, to mitigate the inevitable impacts of entropy and brittleness in AI models over time.
- Creating a persistent and reliable audit trail for all data processing steps to facilitate future analysis, troubleshooting, oversight, and accountability.

We emphasise the critical and ongoing need for model and system maintenance. Unlike "fire-and-forget" solutions, AI systems require consistent care and attention to maintain their effectiveness. This includes:

- Regularly reviewing and updating models to adapt to changing environments and evolving requirements.
- Monitoring the performance of models and addressing any deterioration in accuracy or reliability.
- Implementing robust error handling and feedback mechanisms to ensure that model outputs and potential failures are communicated transparently to users.

We view user interaction with AI outputs as a central feature of the entire system operation, not merely an afterthought. This translates into user-oriented interface considerations that:

- Provide clear and contextual information about the model's outputs, including confidence measures and limitations.
- Offer support mechanisms that augment, and guide informed human decision-making based on AI insights.

Human-Oriented Applications: Our Collaborative Intelligence Initiative

- Prioritising the design of AI systems that complement human expertise and judgement rather than replacing them entirely.
- Recognising the critical role of human oversight and accountability in the deployment and use of AI systems.
- Actively engaging with stakeholders and seeking diverse perspectives to ensure that AI is deployed in a socially responsible and equitable manner.
- We believe in fostering transparency and honesty regarding the trade-offs, limitations, and potential failures inherent in AI systems. This includes:
- Openly communicating the limitations of AI models and the potential for errors or biases in their outputs.
- Providing clear explanations of how AI models are developed and how they reach their conclusions.
- Acknowledging the ethical considerations surrounding AI deployment and engaging in open dialogue with stakeholders about the potential risks and benefits.

By adhering to these principles, we believe our pragmatic approach to AI can lead to the development of reliable, durable, and effective AI tools. These tools are not simply academic musings or fleeting trends, but real-world solutions that address critical challenges across diverse sectors.

Our collaborative work with various clients exemplifies this impactful application of AI. We work closely with them to:

- Understand the specific complexities of their domains and the unique challenges they face.
- Grapple with the legal, policy, and ethical considerations surrounding their desired AI solutions.
- Co-create AI systems that address those complexities on their own terms, ensuring they are contextually appropriate, ethically sound, and operational in real-world settings.
- This approach transcends the realm of performative pronouncements about the potential of AI. It focuses on the practical aspects of building and deploying functional and impactful AI solutions that truly serve the needs of the world around them.

Future Directions and Closing Remarks

As Xaana Pty Ltd concludes this submission to the APH Select Committee on Artificial Intelligence, we would like to outline straightforward recommendations for the federal government that will directly support sovereign AI companies like ours in fulfilling our mission.

We advocate for strategic procurement policies that prioritise Australian-owned AI enterprises for government contracts, especially when it comes to generative AI technology. We urge the federal government to recognise the value of local AI solutions and include sovereign AI companies in the Artificial Intelligence in Government Taskforce and AI and Digital Capability Centers.



We encourage the creation of practical AI governance frameworks with local AI companies as SMEs that include standards for transparency, accountability, ethical usage, and compliance with Australian standards. These frameworks should guide not only public sector deployments but also influence private sector practices, creating a uniform standard across all industries.

We urge government to facilitate the establishment of AI testbeds and living labs where new technologies can be safely tested and refined in controlled, real-world environments. This will accelerate innovation and provide Australian AI companies with valuable data on the performance and societal impact of their technologies.

We recommend government-led and software company driven public awareness campaigns to educate the broader population on the benefits and challenges of AI. This would help demystify AI technologies and promote a more informed dialogue around their use and implications.

With this, Xaana Pty Ltd is deeply committed to advancing AI technology in a manner that upholds our ethos of Collaborative Intelligence, focusing on maximising human capability alongside technological advancements for an AI Future Made in Australia.

We welcome the opportunity to expand on any of our contributions here by appearing before the committee. We look forward to seeing the result of this important inquiry and hearing from you in due course.



Yours faithfully,

Dan Saldi
CEO & Founder of Xaana.AI and Turium AI