



Atlassian's Submission to the Senate Select Committee on Adopting Artificial Intelligence (AI)

Select Committee on Adopting Artificial Intelligence (AI)
Parliament House
Canberra ACT 2600
aicommittee.sen@aph.gov.au

10 May 2024

We appreciate this opportunity to provide feedback to the Senate Select Committee on the opportunities and impacts for Australia arising out of the uptake of AI technologies in Australia, as set out in the Committee's Terms of Reference.

At Atlassian, we build enterprise software products to help teams collaborate, including for software development, project management and content management. As one of Australia's most successful home-grown technology companies — and one that provides products and services to customers around the world — we believe that we are in a unique position to contribute to this consultation.

We know the critical role that technology (including emerging technology) plays in powering the operations of our customers, and the digital economy more broadly. Our customers' technology engineering teams around the world use our products as key tools that help drive their development of technology products and services throughout their lifecycle.

Like many companies, we also leverage the power of machine learning in our own products. For years we have used machine learning to enhance core experiences in our products, from personalised search to recommending people and teams to bring into collaboration. In April 2023, we also announced Atlassian Intelligence, a step forward in bringing the power of AI (including large language models) to our full family of cloud products.¹ We also affirmed our commitment to working with and helping our customers and partners navigate this fast-changing technology landscape responsibly and in line with the values that we all share, in line with our [Responsible Technology Principles](#).²

Since then, as we have continued to build our products with AI-human collaboration at the forefront, we have also continued to learn, grow and share our responsible technology and AI governance practices, including in particular our approach to [responsible technology reviews](#).³

However, as a company that specialises in collaboration and teamwork, we know that individual company principles and commitments aren't enough. We all have a role to play in fostering a thriving, trusted emerging technology ecosystem.

We believe that our regulatory landscape needs to clearly and carefully anticipate what our digital future will mean for Australian organisations and individuals operating in an evolving global economy.

¹ See <https://www.atlassian.com/platform/artificial-intelligence>.

² See <https://www.atlassian.com/trust/responsible-tech-principles>.

³ See <https://www.atlassian.com/blog/strategy/responsible-tech-guide>.

Our approach and proposed model

In late 2020, Atlassian published eight [Principles for Sound Tech Policy](#).⁴ These Principles are intended to not only guide Atlassian's own engagement on important matters of public policy, but to set forth guiding principles for what we believe sound technology-related public policy should look like more broadly..

In line with these Principles, we strongly recommend that the Australian Government and Parliament use this opportunity to design and set forth a purpose-driven, outcomes-focused approach to the governance of AI and other emerging technologies.

Ultimately, we believe that the best response to these emerging technologies is situated within an overarching, coordinated digital regulatory framework that is:

- governed by **core principles**, which may be enshrined in legislation and would set forth a consistent, scalable and risk-based framework for all stakeholders, and inform the formulation and implementation of specific measures and tools within that framework;
- supported through one or more central **advisory bodies** that are capable of providing advice and assisting with coordination and alignment across government agencies and regulators with responsibility across various sectors and areas of law, which would also allow government to build expertise (as to how technology operates, the opportunities and challenges it creates and how best to respond) and connections with industry; and
- bolstered by an appropriate mix of targeted and objective **governance measures, guidance and tools** (including regulatory measures where required), which respond clearly to identified issues in a manner that aligns to and has the benefit of the overarching principles and institutional expertise.

In our view, this proposed model has the flexibility to account for the multi-dimensional nature of many of the opportunities, risks and issues that can be raised by AI. It also allows the governance model to be imbued with Australian values and priorities in its overarching principles, while also retaining the flexibility to both interoperate with and take advantage of emerging international standards and best practice through its agile approach to different policy and regulatory responses.

This recommended overarching model is, of course, only a starting point. In this submission, we set forth our vision not only for *what* framework should be implemented, but also *how* it should be achieved (and *why*) — and how this aligns to work already underway.

Our recommended roadmap: From framework to action

1. Define the playing field

The public release of OpenAI's ChatGPT in late 2022 — and the incredible speed of its adoption by a vast and varied user base —helped to demonstrate the capabilities and the potential of AI to an audience that may not have appreciated just how far AI technologies have already advanced.

But this perception of swift and unprecedented advancement means that just as many individuals and consumers are now discovering the promise and opportunities of AI, they are also discovering the actual and potential issues and harms associated with its use and adoption.

These issues and harms are, in many respects, not new. Instead, the scale and speed involved has led to a perception that we are now treading in territory that is as new and uncharted from a legal and regulatory perspective as it is from a technological one.

⁴ These Principles are also available for download at <https://www.atlassian.com/blog/technology/regulating-technology>.

This is not the case. But we believe that it does illustrate an urgent and critical need to build a common understanding of how the law does — and should — apply to AI and other emerging technologies.

First, we acknowledge that the issues and harms that may be raised by AI can be complex. AI is not a singular technology but an umbrella term for a constellation of technologies, which can involve the use and application of a variety of techniques, processes and tools across a broad spectrum of sectors, contexts and use cases. These range from the productivity solutions that individuals and businesses use on a daily basis (including Atlassian's products) to highly specialised and complex use cases within a range of sectors.

The breadth of these use cases and their application may then give rise to varying, context-dependent risk profiles with legal, regulatory and broader societal dimensions.

Further, we believe that law and regulation, designed carefully, can be a 'force multiplier' for trust: fostering confidence and trust in new industries and technologies, and encouraging their adoption in a way that brings the most benefit to us all.

However, in many contexts, the use of AI is already subject to existing laws and frameworks that apply equally to AI or might be capable of being extended to AI systems under certain conditions. The Human Technology Institute (HTI) has set forth a more detailed overview of the landscape in its May 2023 report on the State of AI Governance in Australia,⁵ and the Australian Government also acknowledged this in its interim response to the safe and responsible AI consultation.

What is not always clear from these outlines and overviews is *why* 'the law' as it stands is not currently being applied, or cannot be applied, to AI. This is because there are a number of different factors that may be at play including, among others:

- where laws do apply, but are not being enforced (for example, due to a lack of awareness or resources) or are insufficient on their own to drive the right outcomes;
- where laws may apply, subject to further clarifications and guidance; and
- where laws could or should apply, with the application of reforms, which may span a spectrum from minor reforms based on first principles (for example, to acknowledge the presence of automated decision-making processes in place of human decision-makers) to major reforms needed to account for actual harms that are not currently addressed by our legal and regulatory frameworks.

Understanding this context is a critical first step to establishing a clear governance model for AI in Australia, as it provides the foundations for understanding the role of law and regulation in responding to emerging technologies like AI — and for implementing a framework that is responsive to the associated issues.

2. Engage with the issue

From the baseline of this common understanding of where we are today and informed by the core principles at the centre of our proposed model, we can then consider how to get to where we want to be.

In particular, the above breakdown of our current landscape can form the basis for prioritising the tools and measures needed to drive towards that outcome. At this relatively early stage, that prioritisation may look like a broad, 'traffic-light' overview of the landscape:

- (1) **Red:** Scenarios of actual, present or imminent harms, for which appropriate interventions don't currently exist and where a proactive approach is justified.
- (2) **Amber:** Scenarios where it is not clear whether appropriate interventions already exist, or whether these may need to be bolstered by further reactive measures. In this case,

⁵ See <https://www.uts.edu.au/human-technology-institute/projects/ai-corporate-governance>.

more long-term work may be needed to identify and implement a program of work and (where relevant) reforms.

- (3) **Green:** Scenarios where laws currently apply or would apply without the need for reform or investigation, but where non-regulatory tools (including greater awareness, resources and guidance) would assist to build and maintain public trust in safe and responsible AI.

We appreciate that it may be difficult to definitively identify which category each use case or scenario may fall into at this stage. This categorisation is accordingly intended to be fluid rather than fixed, and would benefit significantly from the institutional expertise of the advisory body (or bodies) forming part of our proposed model as relevant issues, harms and opportunities emerge and change over time.

3. Treat the ailment, don't kill the patient

Based on the above categorisation and prioritisation, regulatory measures that are targeted towards prevention and remediation of identified harms can then be designed and implemented where needed.

Initially focused on the **red** category, with a view towards addressing **amber** items at later stages, these measures should flow from the principles-based framework and be:

- risk-based by their nature, having evolved from the context of use and application, as well as the categorisation and prioritisation of actual and potential harms involved;
- targeted at addressing where and how the problem or risk under consideration arises, its relationship to the relevant AI technologies and those who provide, implement or use them (as applicable);
- technology-neutral, unless there is a clear and valid justification for implementing a technology-specific measure; and
- with the benefit of the advisory expertise provided above, cognisant of emerging international practice and standards in the circumstances, with a bias towards being interoperable with those standards.

As noted above, while several **red** items may be capable of being identified and responded to in earlier stages, we appreciate that **amber** items are likely to require greater effort to identify and assess, and may still emerge over time as existing concepts and norms are tested by emerging technologies.

Important work has already been done and will continue to be done in this space, including by HTI, the Government's AI Expert Group and the ARC Centre of Excellence for Automated Decision-Making and Society's submission to Government's Safe and Responsible AI consultation⁶. However, we expect that as implications emerge from the widespread adoption of AI and other emerging technologies, this will necessitate ongoing reviews and revision to our legal framework.

4. Tech (and trust) is global

Regulatory and policy options do not have to be limited only to 'hard' law and regulation. The most effective approach is likely to be composed of a broad and flexible mix of targeted 'hard' regulatory measures outlined above, as well as 'soft' law, standards, resourcing and investment, to collectively build a culture of safe and responsible AI in Australia.

In this respect, emerging international responses are illustrative. In particular, many overseas models are demonstrating an increased reliance upon technical and industry standards to drive consistency and uplift industry practice. This includes the US National Institute of Science and Technology's AI Risk Management Framework, the UK AI

⁶ See <https://www.admscentre.org.au/adms-submission-to-the-commonwealth-governments-discussion-paper-on-safe-and-responsible-ai/>

Standards Hub, references to standards and specifications throughout the EU AI Act, and the publication of the International Organisation for Standardisation (ISO) standards relating to AI (including ISO 42001).

These examples are not only useful as learnings for Australia in which measures to consider implementing locally. The inherently global nature of technology (including AI) is such that emerging global standards and best practice can and will have a broader normative effect.

5. Build the foundation for shared success

This mix of approaches will not only be critical to addressing the **green** category above, but also essential to the broader success of the overall model.

In order for Australians (and Australia) to fully and equitably harness the benefits of AI — and be capable of exercising their rights in the digital economy of tomorrow — we will therefore need significant investment in education, skills and awareness.

Case study: To illustrate why this is the case, consider this example:

- In some cases, individuals are not made aware that they are interacting with ‘human-like’ AI systems (like chatbots), or that AI technologies are being used to make certain significant decisions about their lives and livelihoods (like eligibility for insurance).
- Given the potential harms, these scenarios are likely to fall within the **red** category, and new mandatory notification requirements would need to be designed in response.
- However, if individuals receiving that notification aren’t equipped with an understanding of what it means when an AI system is in use, how best to interact with it and how to respond to (and if need be, challenge) its outputs, then those notification requirements could be effectively meaningless.

This investment requires a broad view of lifelong education efforts, with a focus on critical thinking and interdisciplinary skills, aiming to encourage curiosity and discourage tendencies towards unquestioning acceptance and automation bias.

It therefore needs to include education and awareness-raising at the levels of:

- regulators, policymakers and government agencies, with the benefit of the institutional expertise of the advisory body (or bodies) forming part of our proposed model, in order to ensure that our laws can be applied and enforced consistently in the age of AI;
- the workforce, to encourage responsible governance of AI throughout all levels of organisations (including in line with the recommendations set forth in HTI’s report on the State of AI Governance in Australia); and
- our education and training system, so that we raise the awareness of AI use cases, that these skills are embedded as early as possible and that young Australians are well-equipped for our digital future.

Our support for the Australian Government’s current approach

In line with our proposed model above, we believe the approach outlined in the Australian Government’s interim response to the safe and responsible AI consultation reflects a welcome, and necessary, first step towards implementing a best practice framework.

In our view, the interim response acknowledges the current shifting AI landscape and the associated open questions, unknowns and potential limitations, as well as opportunities from the adoption of this emerging technology. It acknowledges the role of both current and newly proposed laws in responding to the issues raised by AI, the need to investigate regulatory and non-regulatory mechanisms, and the ways in which international developments and frameworks can influence domestic best practice.

It also identifies, and balances, both immediate and longer-term priorities in a way that establishes the foundations for a flexible framework that can evolve alongside this shifting landscape.

We look forward to the advice provided by the AI Expert Group on the definition of "high-risk" use cases and the mechanisms for regulating the use of AI in these settings.

Atlassian has also been engaging with the process lead by the National AI Centre to develop voluntary AI safety standards and guidance to improve the testing, transparency and accountability of AI systems being developed and deployed in Australia. As a proud advisory partner of HTI, we have also continued to partner with HTI on its program relating to The Future of AI Regulation in Australia, which we believe will also be of significant benefit to the Government's ongoing work in this area.

Atlassian would be pleased to further discuss these comments with the Committee at its convenience.

Yours sincerely,

Anna Jaffe
Director of Regulatory Affairs & Ethics
Atlassian

David Masters
Head of Global Policy & Regulatory Affairs
Atlassian