Comments on Adopting Artificial Intelligence (AI)

29 April 2024

Workday appreciates the opportunity to provide comments to the Select Committee on Adopting Artificial Intelligence (AI). Workday is a leading provider of enterprise cloud applications for finance and human resources, helping customers adapt and thrive in a changing world. Workday applications for financial management, human resources, planning, spend management, and analytics are built with artificial intelligence and machine learning at the core to help organisations around the world embrace the future of work. Workday is used by more than 10,000 organisations around the world and across industries – from medium-sized businesses to more than 50% of the *Fortune* 500. With offices in Brisbane, Melbourne, and North Sydney, two in-country data centres, and a customer support presence, we are proud of our robust offerings in Australia. Workday serves major Australian customers including Atlassian, Canva, the Commonwealth Bank of Australia, Latitude Financial Services, One Rail Australia, QANTAS, Reece Group, Telstra, and St Vincent's Health Australia.

At Workday, we believe that artificial intelligence ("AI") is powering the future of work by unlocking human potential, driving business value, and enabling our customers and their employees to focus on more strategic and fulfilling work. Consistent with our commitment to ethical AI, Workday has been helping to lay the groundwork for smart AI safeguards since 2019. Building on our call for AI regulation, Workday has taken a leading role in AI policy discussions at the federal, state, and local level in the United States. We have also partnered with officials in the European Union on the Artificial Intelligence Act, as well as governments around the world, including the United Kingdom, Singapore and Japan, to provide thoughtful and concrete policy approaches to responsible AI.

Workday supports the development of safe and responsible AI policies, regulations and practices, and offers the following comments that will help advance Australia's ongoing efforts to promote safe and responsible AI. Please do not hesitate to contact Eunice Lim, Director Corporate Affairs – AP&J, at eunice.lim@workday.com if you have any questions or would like further information.

I.   AI and a Skills-based Workforce

Workday believes that AI can positively transform how people and organisations work. The AI tools we deliver to our customers help their employees make more informed decisions by surfacing new insights, identifying opportunities for career development, and improving their day-to-day work experience by simplifying labour-intensive tasks. Our guiding principle is that AI should be used in ways that augment, rather than displace people.

While employers understand the benefits of using AI to address both labour-intensive and time-intensive work, they wrestle with growing skills gaps in AI, AI-adjacent, and more traditional roles. A key question then is how such a skills gap should be addressed while also equipping workers and employers with the resources to navigate the coming changes in the workplace.

At Workday, we believe that a shift to a skills-based approach to talent is the best way forward. A skills-based approach to talent refers to an emphasis on what a person can do or learn, rather than solely on their credentials. It can increase both workers' and employers' ability to respond more nimbly to shifts in the economy while expanding opportunities for diverse and underrepresented communities.

AI is essential in enabling a skills-based approach to the workforce at scale. Skills data is often complex, noisy, and voluminous often with the same terms used to describe multiple skills and multiple words used to convey the same skill. Deriving actionable insights has historically been a hurdle. AI's unique ability to parse massive amounts of skills data generate insights into existing and in-demand worker skills which drive informed decision-making in the current and future job market. AI also maximises the impact of workforce development efforts by ensuring that policy interventions are ready-made to tackle these challenges. Workday supports a full consideration of not only the impact technology can have, but also the opportunities that AI tools can unlock when applied to skills-based approaches to talent.

II.   Responsible AI Governance Builds Trust

While AI presents incredible opportunities to unlock human potential, there is also the risk of unintended consequences. In a recent research conducted by Workday, Closing the AI Trust Gap, it was found that the AI trust gap in Australia is worryingly large and Australians' scepticism of AI is higher than the global average. Sixty percent of Australians are worried about the trustworthiness of AI, the highest among all countries surveyed, and only half (51%) are confident their organisation can clearly say AI will not improve work.

As a cloud-native enterprise software company, Workday learned early on that rigorous investments in technology governance are critical to earning and retaining our customers' trust. This is why Workday has put in place a robust Responsible AI program and supports regulatory and policy measures incentivising organisations to establish internal AI risk management programs. Workday's Responsible AI program includes:

- **Leadership Commitment** from a Responsible AI Advisory Board that is led by our General Counsel and counts our Chief Compliance Officer, Chief Technology Officer, and Chief Diversity Officer among its members.
- **Dedicated Resources** that include a team of social and data scientists and technology experts that report to our Board of Directors through our Chief Compliance Officer and that develops and maintains Workday's responsible AI governance framework. The team receives cross-company support, including from responsible AI champions who provide subject matter expertise so that AI products are developed in accordance with Workday's AI principles.
- **Responsible AI Guidelines and Review Processes** that operationalise our principles through AI development guardrails, turning them into documented practices and assessments. Our product development teams use tools to evaluate a potential AI feature's risk profile before we write any code. AI tools intended for use in consequential decisions, such as hiring or promotion, are treated as high-risk.
- **Disclosure** to equip our customers with a clear understanding of how our AI tools are developed

and assessed, as well as transparency and choice in how their data is used.

Although Workday has taken these steps to develop AI in a responsible manner, we recognise that the lack of public trust in AI must be addressed across industry and also complemented by regulatory guidance from policymakers.

Workday's trust gap report also revealed that the Australian workforce views effective regulation of AI and data as a step towards increasing the trustworthiness of AI.

III.   Risk-based Approach to AI Governance

The Australian Government has been making strides to address the regulatory gaps in AI through the work of the Department of Industry, Science and Resources (DISR). At Workday, we strongly support smart regulations on AI that support innovation and build trust, including AI tools used for consequential decisions like hiring, promotion, and termination. We have been engaging and providing DISR with our thoughts on a path forward for AI governance and meaningful safeguards. In fact, DISR's recently published interim response to the Safe and Responsible AI consultation echoed themes consistent with Workday's AI policy positions, including the need to take a risk-based approach to AI regulation.

A risk-based approach to AI governance means applying rules to contexts where AI carries the highest risk of potential harm to individuals. As AI systems are and will be used in a wide variety of scenarios, policies and regulations must be designed to reflect this diversity in risk profiles. A risk-based approach towards AI governance would allow for the benefits of AI technologies to be harnessed while providing the necessary safeguards to minimise the potential harms arising from the use of AI. Many countries are recognising the merits of, and converging on, a risk-based approach to AI governance, including the European Union, Canada, the United Kingdom, and Singapore.

To achieve a risk-based approach to AI governance in practice, Workday recommends the following:

a)   Scope consequential decisions based on the risk of harm.
     The AI ecosystem is broad, encompassing a diverse range of technologies and use cases and a wide array of stakeholders. Because the risks of AI vary according to its use-case, any regulatory measure should focus on specific applications of the technology that pose higher risks to the public, yet should be flexible enough to account for the unique considerations that may be implicated by specific use cases. The risk profile of specific use-cases will also vary within any given sector and demands a nuanced approach. With workability as a key goal, it is important to define the scope of impacts that are the focus of any definition of what is and is not high-risk and avoid assigning risks solely based on a sectoral approach.

b)   Consider the level of automation associated with AI use-cases.
     In addition to the scope of any high-risk AI definition, regulations should also consider the level of automation associated with the AI tool. For example, fully automated AI tools that seek to replace human decisions rather than those that augment human decision-making inherently pose a higher level of risk than those that assist a human and in which the human remains in control of

any decision-making. Truly automated tools can make consequential decisions at a higher volume and velocity than a human and lack accountability. By contrast, when an informed human is in the loop, they can leverage an AI tool's insights and remain in control of, and accountable for, the final decision.

c) <u>Distinguish between Internet-Trained Large Language Models (LLM) and domain-specific models</u>.

The commercial LLMs used in generative AI are often trained on data scraped from the Internet, i.e. Internet Trained LLMs (IT LLMs). These IT LLMs are used across a wide variety of applications and domains, and often have uncertain provenance associated with the data on which they are trained. With companies moving to leverage generative AI for narrow, domain-specific applications trained on curated and high-quality data sets with clear ownership, it is important to recognise that IT LLMs and domain-specific models present different levels of risk profile. A risk-based approach towards AI governance would take into consideration the variation in risk profile between IT LLMs and domain-specific models while also ensuring sufficiently rigorous transparency requirements are applied to IT LLMs to ensure it is possible to leverage them responsibly.