



Australian
National
University

ChatGPT

Capabilities

Remembers what user said
earlier in the conversation

Allows user to provide follow-
up corrections

Trained to decline inappropriate
requests

Limitations

May occasionally generate
incorrect information

May occasionally generate
harmful instructions or content

AI Adoption in Australia

A Partnership between Google and ANU

Authors

From ANU School of Cybernetics:

Jessica Herrington, Safiya Okai-Ugbaje, Ellen O'Brien

From ANU School of Politics and International Relations:

Nicholas Biddle



Australian
National
University

AI Adoption in Australia

A Partnership between Google and ANU

Authors: Jessica Herrington,¹ Nicholas Biddle², Safiya Okai-Ugbaje ¹ Ellen O'Brien¹

1. School of Cybernetics, Australian National University
2. School of Politics and International Relations, Australian National University

May 2026

*Contact: Jessica.Herrington@anu.edu.au ; or Nicholas.Biddle@anu.edu.au

School of Cybernetics
ANU College of Systems and Society
+61 2 6125 8121

E: cybernetics@anu.edu.au

W: cybernetics.anu.edu.au

The Australian National University

Canberra ACT 2601 Australia

www.anu.edu.au

TEQSA Provider ID: PRV12002 (Australian University) | CRICOS Provider Code: 00120C

Executive Summary

Generative Artificial Intelligence (GAI) has moved rapidly from a niche technological development to a widely recognised and increasingly used tool within Australian society. This report provides one of the most comprehensive examinations to date of Australians' familiarity with, use of, and attitudes toward GAI, as well as their expectations for governance and acceptable use. It draws on a nationally representative survey of Australian adults conducted in late 2025, supplemented by targeted interviews with Australian organisations adopting or preparing to adopt GAI systems. Together, these data provide new insights into how Australians understand and engage with GAI, who is benefiting from its emergence, and where capability, trust, and governance gaps are likely to shape future adoption and policy needs.

Methodology

This research is composed of two parts: a survey component, and semi-structured interviews. The survey component of this research used a weighted, nationally representative sample of Australian adults. Respondents were asked about their general interest in emerging technologies; their familiarity with Generative AI; their use of GAI across home, work, and study settings; their perceived personal benefits and concerns; their views on governance and responsibility; and their reactions to experimental vignettes testing the acceptability of GAI in specific contexts. Survey estimates were weighted to population benchmarks to ensure national representativeness. The semi-structured interviews were conducted with a range of organisations across sectors, focusing on adoption drivers, workforce impacts, regulatory expectations, and organisational readiness. Findings from the survey and interviews were analysed separately and then synthesised to support broader conclusions and implications for policy and practice.

Key Findings from the Survey

- Awareness and interest in emerging technologies are high, with more than half of Australians expressing interest in technologies such as GAI. However, familiarity and attitudes are strongly patterned by age, education, language background, and place of residence.
- Understanding of GAI is widespread but uneven, with nearly three-quarters reporting at least a basic or moderate understanding, but only a small minority reporting strong familiarity. Younger and more educated Australians consistently report higher levels of understanding.
- Use of GAI is already common, with almost half of Australians having used GAI at least once. Adoption is much higher among younger Australians, those with tertiary qualifications, and people living in capital cities — indicating emerging digital capability divides.
- Users report clear individual-level benefits, especially time savings. Many use that time to undertake more complex tasks, handle greater volumes of work, or engage in more creative activities. Students and workers who use GAI report meaningful improvements in study quality and job performance.
- Concerns about data, skills, and critical-thinking capability persist, with majorities identifying lack of training or skills, privacy concerns, or fears of over-reliance as limiting their use.
- Australians support shared responsibility for governing GAI, with a particularly strong role for government. Majorities support labelling of GAI-generated content and express high levels of concern about political uses of GAI.
- Acceptability of GAI use is sensitive to context, with vignette experiments showing that support decreases when a greater share of content is produced by AI. Individual characteristics — such as gender, age, and education — also shape attitudes, with women and older Australians generally more cautious.

Key Findings from Organisational Interviews

- Organisations see AI as a transformational technology with potential to reshape workflows, customer service, knowledge work, and internal processes.
- Many organisations are in early adoption phases, focusing on testing and experimentation, with a strong emphasis on risk identification, change management, and staff training.
- Governance complexity is a central theme: organisations want clear standards, shared responsibilities, and greater transparency from technology providers.
- Different levels of access to GAI tools and governance frameworks at the organisational level within the Australian Public Service. While recognising different risk levels, there was consensus for standardisation of access and clearer regulatory direction.
- Internal capability gaps—spanning digital skills, leadership confidence, and organisational readiness—are seen as major determinants of successful adoption.

Implications

Taken together, the findings in this report indicate that Generative AI is already embedded in the lives of many Australians, but its benefits are not yet equally distributed. Younger, more educated and metropolitan Australians are at the forefront of adoption, while older Australians and those with lower levels of education report lower use, weaker understanding, and greater concerns. As GAI becomes increasingly integrated into everyday products, services, workplaces, and public institutions, these gaps risk becoming barriers to equitable opportunity, workforce participation, and trust in digital government and democratic processes.

The survey and interview evidence therefore points to a clear need for targeted capability-building, transparent governance frameworks, and consistent labelling and disclosure practices. Organisations and government agencies will need to invest in training, develop clear internal and external standards for GAI use, and address privacy, data, and critical-thinking concerns through well-designed communication and safeguards. Ensuring that the benefits of AI are realised across the whole population—and that its risks are managed in a way that sustains public trust—will require coordinated action across government, industry, and civil society.

Key Policy Insights

Invest in national AI capability building and digital inclusion

Strengthen educational pathways, ensure equitable access across states and regions.

Develop standards and education for trustworthy AI

Prioritise transparency, model reliability, explainability safeguards, as well as guidance and education on appropriate use.

Support sector-specific AI adaptation and fine-tuning capability

Enable government, industry, and education sectors to more effectively tailor existing AI models for Australian contexts rather than over-reliance on generic global tools.

Embed cultural and ethical considerations into national AI governance

Integrate First Nations data sovereignty principles, and values-led design into AI policy.

Reinforce human elements

Policies should recognise that confidence and capability depend on human factors such as identity, narratives, and trust as much as on technical infrastructure.

Position Australia as a leader in sustainable, values-based AI

Leaders can leverage regulatory strengths, address sustainability concerns, and facilitate cross-sector collaboration to shape a distinctive national approach.

Contents

Executive Summary	iii
Methodology	iii
Key Findings from the Survey	iii
Key Findings from Organisational Interviews	iv
Implications.....	iv
Key Policy Insights.....	iv
Contents	v
1 Introduction.....	1
1.1 Our Approach	2
2 Surveys.....	4
2.1 Data.....	4
2.2 Generative AI familiarity and views.....	4
2.3 Use of Generative AI.....	7
2.4 Perceived impact of Generative AI	10
2.5 Generative AI and governance	16
2.6 Experiments in acceptable use of Generative AI	21
3 Interviews.....	25
3.1 Tool Use and Practices.....	25
3.2 Motivations and Drivers of Adoption	27
3.3 Challenges & Inhibitors.....	29
3.4 Emergent Narratives.....	31
3.5 The Future of AI Adoption: A Uniquely Australian Perspective	32
4 Summary and Conclusions: What's Next.....	35
5 References	37



1 Introduction

One of the challenges of Artificial Intelligence (AI) is that the term holds many definitions – from rules-based algorithms; to machine learning; to stochastic algorithms; to agentic AI; to speculative machines that possess general intelligence, and many distinctions in between. In reality, AI is a constellation of many technologies – and while the ways in which AI turns up in the public vernacular are not always definitionally precise, they are often helpful to understand how the public presently encounters AI. Currently, AI is synonymous with Generative AI (GAI).

Since the transformer model revolutionised AI research and development in 2017, the public has increasingly encountered AI in the form of large language models and multi-modal models. Now, through the release of generative AI tools in the 2020s, public perception of AI is overwhelmingly informed by these generative applications of AI.

Indeed, GAI has seen rapid and remarkable uptake. The pace and scale of GAI uptake outpace both personal computers (PCs) and the Internet in their early stages of adoption.¹ Just this year, findings from the McKinsey 2025 report showed continued increases in GAI use, with 72% of survey respondents reporting frequent use of GAI². This is a jump from 65% in 2024, and 55% in 2023³.

This continued growth in GAI usage signals a shift in how people interact with AI. We have moved from a market where AI-powered tools have mostly served as enterprise applications, or been opaquely embedded into products, to a market where GAI-powered tools are widely and deliberately adopted for work, home, and study. Seemingly, AI is becoming quotidian and increasingly domesticated through GAI.

As researchers, we have an opportunity to study the levers that are determining this transition to AI and particularly GAI at scale. For example, does a conversational mode of interaction support people's receptiveness to GAI-powered tools? Might it shape the role that GAI plays in a person's life? Does GAI's use in the domestic spaces start to break down people's historic aversion to trusting AI? How is GAI augmenting our skills? How is this impacting our workplaces? Our industries? And, what will the impact of GAI at scale look like? Will ubiquitous GAI reshape our information landscape? Will it change the public perception of AI writ large?

In this study, we intend to look at a constellation of factors that are shaping AI adoption by individuals and enterprises. We are focusing on the Australian context and plan to draw cross-country comparisons with published studies in other jurisdictions. Our research spans both individual-level and enterprise-level adoption. We seek to examine what the true use of AI in Australia is, and what impacts of this use are beginning to emerge.

¹ Bick, Alexander, Adam Blandin, and David J. Deming. "The Rapid Adoption of Generative AI." Working Paper. Working Paper Series. National Bureau of Economic Research, September 2024. <https://doi.org/10.3386/w32966>.

² "The State of AI in Early 2025 | McKinsey." Accessed March 23 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

³ "The State of AI in Early 2024 | McKinsey." Accessed January 15, 2025. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

1.1 Our Approach

In this study, we seek to understand AI adoption in Australia across multiple levels, including adoption by individuals, adoption by organisations, and adoption within key sectors to inform future policymaking. To conduct this research, we have undertaken a mixed-methods approach that includes a population survey, a survey of organisations, and a series of expert interviews to represent a variety of sectors and demographics within Australia.

In our population survey, we have defined GAI as ***deep-learning models that can generate high-quality text, images, and other content in response to the user's prompt, based on data they were trained on.*** This definition includes models such as ChatGPT, Google Gemini, Dall-E, Claude and several others.

To address AI adoption more broadly (including AI applications beyond generative AI), in the organisational survey we are targeting responses based on a definition of AI as, *as a machine-based system that, for explicit or implicit objectives, infers from the input it receives how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments.*

We make this point of difference between the population and organisational surveys intentionally. At the organisational level AI applications are often strategically procured to meet a range of objectives that are not all covered by generative AI – unlike at the individual level where AI applications other than GAI may simply be embedded in already established applications. We anticipate useful insights from comparing across AI applications against other factors that affect adoption in organisations.

Goals for this research, include:

- ***Understanding the demography of (G)AI adoption.***

Are there differences in adoption rates across age, gender, linguistic or cultural background, socioeconomic status or level of education that could inform future policymaking?

- ***Understanding AI adoption by institutions.***

Which types of organisations are adopting AI? And, are there patterns in industry, geography (e.g. urban / regional or state to state), reach (e.g. local, domestic or multinational), or organisational scale that are shaping adoption behaviours?

- ***Understanding the rationale for AI adoption and its impact.***

Can we learn anything from why early adopters have decided to use AI? What benefits will these changes bring (e.g. productivity, better work life balance, transformative capability change, occupational health and safety)? And, are there emerging practical effects of AI adoption we could seek to investigate in future?

For our analysis, we are focused on identifying multiple factors that drive certain behaviours and systemic outcomes related to AI adoption. By taking a multi-level approach (individual, organisation and sectoral) to unpacking adoption, our study will cut across both the individual drivers, organisational enablers, and environmental conditions of AI adoption to unpack the systemic factors at play in AI adoption and emerging impacts of adoption trends. To undertake this approach, we are looking at a range of factors, including signals of what the impacts of AI at scale may be. As an outcome of this approach, we anticipate that our research will provide a novel offering among the emerging literature on GAI and AI adoption, which has tended to take a narrower scope of analysis to-date, and does not look closely at how actors across these multiple levels of a broader system might be shaping adoption behaviour.

Our approach to this study has been inspired by past work that explores the drivers and impacts of transformational technologies at inflection points as these technologies scale on a trajectory to ubiquity. For example, a study conducted on the SIP (Singapore internet project) – a project to connect the whole island of Singapore to the internet – in collaboration with the WIP (World internet project) – conducted a longitudinal study of a 1000-person sample representative of the

Singaporean population, published in 2002⁴. This study looks beyond usage patterns of the Internet to study factors such as impacts to lifestyle that are shaped by Internet use; effects to perceptions of value and trust in various information sources; impacts to trust in institutions during broad uptake of the Internet; attitudes towards E-commerce; and attitudes to policy and regulation linked to the Internet. Outcomes of this study helped to explain how specific affordances of the Internet, including access to information at an unprecedented scale, and communication via new modes, had implications beyond adoption, and were shaping domestic habits, social relationships, and political perspectives on Internet adoption in Singapore.

As GAI is already widely adopted, studies such as the Singapore Internet study offer a blueprint for thinking beyond only the drivers of adoption, to emerging impacts and implications of adoption that will be an important feature for understanding the future of AI adoption, and a key element of our study.

⁴ Kuo, Eddie C. Y., ed. *Internet in Singapore: A Study on Usage and Impact*. Singapore: Times Academic Press, 2002.

2 Surveys

2.1 Data

Individual-level data for this project was collected as part of the ANUpoll series of surveys. Data was collected in partnership with the Online Research Unit (ORU) through the Australian Consumer Panel. Data collection commenced on the 8th of October with a pilot sample. Full data collection commenced the following day, with data collection finishing on the 19th of October with 3,577 respondents. All survey data in this report is weighted to the Australian adult population.⁵

2.2 Generative AI familiarity and views

The first question asked on the survey that was specific to AI readiness was ‘Firstly, a general question about your attitudes towards technology. How interested would you say you are in emerging technology, for example generative AI, quantum computing, autonomous vehicles, biotechnology?’ While this is more general than Artificial Intelligence, it has been shown in the literature to be highly predictive of attitudes related to specific technology. For example, individuals with a generally positive stance toward technology are more likely to adopt and support health-related mobile applications, perceiving them as extensions of progress rather than sources of risk⁶. Conversely, those with technophobic tendencies tend to exhibit scepticism or resistance toward even well-established technologies, driven by underlying beliefs about control, privacy, or societal disruption.⁷

The majority of Australians – 52.7 per cent – are either interested (37.9 per cent) or very interested (14.8 per cent) in emerging technology. We analysed the predictors of this interest. Using the ordered probit model, we estimated the demographic, socioeconomic, and geographic predictors, with higher values in the model associated with a greater interest. Figure 2.1 gives descriptive statistics of the key explanatory variables that we found to be important (that is without controlling for other characteristics)

Gender was strongly associated with interest in technology. 61.5 per cent of males say that they are interested/very interested compared to 44.2 per cent of females. Age is also associated, with the biggest gap between those aged 18 to 54 years (59.1 per cent interested/very interested) and those aged 55 years and over (44.2 per cent).


The socio-demographic variable with the strongest association was education. Only 35.2 per cent of Australians that have not completed Year 12 said they were interested/very interested in emerging technologies. This rises to 48.8 per cent for those that have completed Year 12 or have a non-degree qualification, and even more so to 71.3 per cent for those with a degree.

⁵ The data has been weighted using iterative proportional fitting (raking) using four benchmark variables. The first two – age and sex – come from ABS population estimates. Education levels come from the 2021 Census. And current employment comes from the Labour Force Survey. Except where stated, any standard errors that are calculated are based on bootstrap standard errors with 250 replications.

⁶ Xie, Z., Or, C. and Ye, X.C., 2026. Acceptance of mobile health applications by the general public: the roles of technology acceptance model constructs, health-app-specific factors, and socio-demographic moderators. *Applied Ergonomics*, 130, p.104655.

⁷ Novozhilova, E., Vu, C. and Katz, J., 2025. From moral panic to normalization: comparing users and non-users of AI companionship apps. *AI & SOCIETY*, pp.1-19.

Figure 2.1 Interest in emerging technology, by gender, age, and education



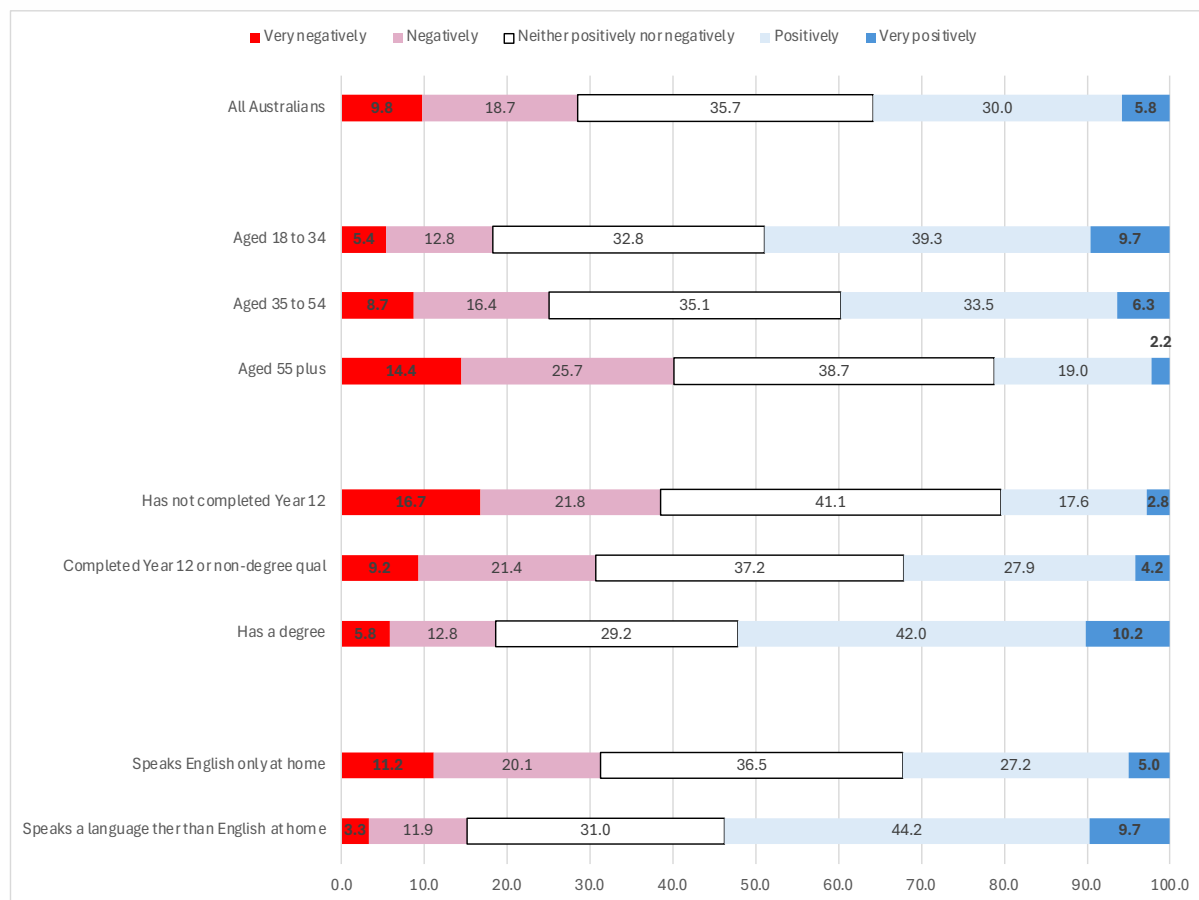
We next gave respondents a generic description of Generative AI as follows: For this survey we are defining Generative Artificial Intelligence (AI) as **deep-learning models that can generate high-quality text, images, and other content** in response to the user's prompt, based on data they were trained on. Examples of these models include, ChatGPT, Dall-E, Claude and several others.' [bold in the original]

When asked for their views based on this description, there was a slightly higher proportion of Australians that viewed Generative AI positively (30.0 per cent positively and 5.8 per cent very positively) than viewed it negatively (9.8 per cent very negatively and 18.7 per cent negatively). However, the modal or most common response was neutral, with 35.7 per cent of Australians saying they viewed it neither positively nor negatively.

There were a number of factors that were associated with the positivity of a person's view towards AI when we estimate using an ordered probit model. We find that gender (females less positive), location (those outside of capital cities less positive) and country of birth (those born overseas in a non-English speaking country less positive) all have a significant association. However, these differences do not remain when we control for interest in technology (which has a strong and very positive association). Three variables that do remain significant even after controlling for interest in technology are age, education, and language spoken at home

Younger Australians are far more positive towards Generative AI (49.0 per cent positive or very positive for those aged 18 to 34) than those in the middle part of the age distribution (39.8 per cent for those aged 35 to 54) or those aged 55 years or over (21.2 per cent). Those with a degree are far more positive (52.2 per cent) than those that have either completed Year 12 or a non-degree qualification (32.1 per cent) and those that have not completed Year 12 or a qualification (20.4 per cent). Those that speak a language other than English at home (53.8 per cent) are more positive than those that speak English only (32.2 per cent).

Figure 2.2 Positive or negative views towards Generative AI, by age, education, and language spoken at home



Respondents were next asked about their level of understanding of generative AI, ranging from saying that they understand Generative AI very well (given by 6.6 per cent of respondents) to never having heard of it prior to the survey (3.5 per cent). In between these two extremes, the most common response was that a person felt they had a basic understanding of Generative AI. This was given by 45.7 per cent of respondents, with a further 24.6 per cent saying they have a moderate understanding and 19.6 per cent saying they do not understand it, but have heard of it.

There are some differences in the perceived level of understanding of Generative AI, particularly when we compare those that feel they understand it moderately or very well, compared to those that feel they have a basic understanding or less. Males are more likely to say they have at least a moderate understanding of Generative AI (34.3 per cent) than females (27.8 per cent). Younger Australians are far more likely to say they have a moderate or greater understanding than older Australians, with a reasonably linear decline across the age distribution (from 48.2 per cent of those aged 18 to 24 to 10.1 per cent of those aged 75 plus). Those with a degree are far more likely to say they have at least a moderate understanding (43.5 per cent) than those that have completed Year 12 only or have a non-degree qualification (29.2 per cent) and in particular those that have not completed Year 12 or a qualification (17.4 per cent).

The final question that we ask on familiarity with Generative AI is when a person says they first heard of it. We only ask the question of those that had heard of it prior to the survey itself.

- 22.4 per cent of respondents say they heard of Generative AI 3 years or more ago
- 29.5 per cent of respondents say they heard of it 2 years ago
- 29.7 per cent of respondents say they heard of it 1 years ago
- 9.6 per cent of respondents say they heard of it 6 months ago
- 8.8 per cent of respondents say they heard of it less than 6 months ago

Summary and implications

Overall, Australians report a high degree of awareness of generative AI, broad interest in emerging technologies, and substantial variation across demographic groups. More than half of Australians say they are interested in emerging technologies, and this general technological orientation is a strong predictor of attitudes toward generative AI. Views towards Generative AI are more positive than negative on balance, although neutrality remains the modal response, suggesting that many Australians are still forming opinions.

Younger Australians, those with a university degree, and those who speak a language other than English at home tend to hold more favourable views, even when controlling for their broader interest in technology. Self-assessed understanding follows a similar pattern: while only a small proportion feel they understand generative AI very well, almost three-quarters report at least a basic or moderate understanding, with much higher levels among younger and more educated respondents.

The uneven distribution of interest and understanding points to the possibility of emerging disparities in who is positioned to benefit from Generative AI, highlighting the importance of inclusive strategies that support engagement across all demographic groups.

2.3 Use of Generative AI

Level and frequency of use

The next few questions in the survey ask about the use of Generative AI. Before asking the question, we tell respondents 'We are interested in understanding your patterns of Generative AI use. When we talk about use we are interested in your deliberate use of Generative AI, rather than your use of tools or products that may incidentally have Generative AI embedded into it.' Based on this definition, 48.6 per cent of Australians say that they have used Generative AI.

The predictors of the use of Generative AI are very similar to the predictors of awareness and understanding. But there are some differences. Gender does not have an association with whether or not a person says they have used generative AI. Age has a very strong association. More than seven-in-ten Australians aged under the age of 35 said they have used Generative AI, compared to 17.2 per cent of those aged 65 to 74 years, or 12.3 per cent of those aged 75 years and over.

Education also has an association. Those with a degree are far more likely to say they have used Generative AI (66.7 per cent) than those that have completed Year 12 only or have a non-degree qualification (47.3 per cent) and in particular those that have not completed Year 12 (24.5 per cent). Those that live outside of a capital city are far less likely to have used Generative AI (39.5 per cent) compared to those living in a capital city (53.2 per cent)

Those that have used Generative AI are far more positive about it than those that have not. Specifically, 54.7 per cent of those that have used Generative AI have a positive or very positive view, compared to 18.7 per cent of those that have not used generative AI. These differences are also present when we re-estimate the ordered probit model from above with use of AI as an additional independent variable

For those that have used Generative AI, we asked when they first started using it.

- 9.1 per cent of respondents say first started using Generative AI 3 years or more ago
- 23.5 per cent of respondents say it was 2 years ago
- 35.1 per cent of respondents say it was 1 years ago
- 18.8 per cent of respondents say it was 6 months ago
- 13.5 per cent of respondents say it was less than 6 months ago

For those that have used Generative AI, we also asked how frequently they use Generative AI at work (for those that are employed), as a student (for students), and at home. For users who are working, 9.8 per cent use Generative AI everyday at work, 23.4 per cent use it most days, 24.3 per cent use it a day or two per week, and 42.5 per cent use it less than once per week. For users who are studying, 13.6 per cent use Generative AI everyday as a student, 28.6 per cent use it most days, 28.6 per cent use it a day or two per week, and 29.2 per cent use it less than once per week.

For all users, 8.6 per cent use Generative AI everyday, 18.9 per cent use it most days, 24.3 per cent use it a day or two per week, and 48.2 per cent use it less than once per week.

For those that have used AI, we asked ‘Over the next 12 months, to what extent do you expect your use of Generative AI to increase or decrease?’. Far more respondents think that their usage will increase rather than decrease. 17.9 per cent said that it would increase substantially, with the most common response simply being that it will increase (42.0 per cent of respondents). Only a very small per cent of respondents think their usage will decrease (2.0 per cent decrease, 2.7 per cent decrease substantially), with a sizable minority (35.4 per cent) saying that their usage of Generative AI will stay about the same.

For those that had not used Generative AI, we asked ‘Over the next 12 months, to what extent do you expect to use Generative AI?’ Almost two-thirds of respondents (64.5 per cent) said that they would continue to not use it all. A little over one third (34.1 per cent) said that they would use it a little, with only 1.4 per cent of those that had not used Generative AI saying that they would use it a lot.

Still focusing on those that have not used Generative AI, those that nonetheless view it positively are far more likely to say they will commence using Generative AI than those that view it negatively (75.8 per cent compared to 26.5 per cent).

We asked respondents about the frequency of specific activities that they used Generative AI for. We asked about twelve activities ‘at work, as a student, or looking for work’ (Figure 2.3), and a further eight activities at home (Figure 2.4). These questions were asked only of those that have used Generative AI, with five possible frequencies, and as sixth response option for those that do not undertake those tasks.

The most frequent activity that Generative AI is used for in both a professional and home setting is searching for information and facts. This is followed by writing communications

Figure 2.3 Frequency of use of Generative AI for activities at work, as a student, or looking for work





Figure 2.4 Frequency of use of Generative AI at home



How Generative AI is used

Respondents that had used Generative AI were asked to choose one of six analogies to describe their interactions. They were also offered an 'other' category, with the availability of free text input to describe their interactions.

The most common response given is that respondents use Generative AI as an information resource. This was the analogy chosen by 47.1 per cent of those that have used Generative AI. Other analogies that a sizable minority of people reported were personal assistant (21.2 per cent) and collaborator (10.8 per cent). Less frequently given responses were friend or confidante (6.9 per cent), advisor (5.8 per cent) or coach (3.9 per cent).

An additional 4.4 per cent of respondents gave 'other' as their response. The free-text that respondents gave to this question surfaced some of the more negative attitudes towards it. For example, some responses include: An unpleasant necessity these days, Annoyance, Enemy (given by eight respondents), Nemesis (given by two), and That friend that knows stuff that I don't like.

Respondents were also asked whether they used Generative AI in a conversational way. Of those who have used Generative AI, 39.4 per cent say that they use it in a conversational way.

Following the question on conversational approach to Generative AI, respondents are reminded that 'Generative AI is often embedded into products, platforms and services. For example, Generative AI can be used in search engine features that help to summarise search results.' They are then asked 'Would you say you are aware of when you are interacting with Generative AI that is embedded into a product, platform or service?' This was asked of all respondents, not just those that have used Generative AI. 54.8 per cent of Australians are aware when they are interacting with embedded Generative AI.

Awareness of the often-embedded nature of Generative AI is much higher for those that have used Generative AI (72.8 per cent) than those that have not (37.8 per cent).

Even controlling for whether or not someone has used Generative AI, there are still differences in this awareness across demographic and socioeconomic variables in the dataset. Younger Australians are more likely to be aware that Generative AI is embedded. The biggest difference is between those aged 18 to 34 years (67.9 per cent were aware) and those aged 35 years and over (49.1 per cent were aware). There are also differences by education with those that have not completed Year 12 or a qualification having a much lower awareness (38.3 per cent) than those that have completed Year 12 only or have a non-degree qualification (55.6 per cent) and particularly those that have a degree 64.8 per cent

Summary and implications

Australians' reported use of generative AI is substantial, with nearly half of the population (48.6 per cent) having used generative AI at least once. Consistent with earlier findings on familiarity and interest, use is strongly patterned by age and education. More than seven in ten Australians under 35 report using generative AI, compared with less than one in five of those aged 65 or over, and two-thirds of degree-holders report use compared with one-quarter of those who have not completed Year 12. Use is also higher among residents of capital cities than those living outside them. Frequency of use varies by context: around one in ten employed users report using generative AI every day at work, and around one in five use it most days, with similar patterns among students. Across all settings, the most common activities involve searching for information, writing communications, and undertaking content-related tasks—consistent with the affordances of current large language models.

For many users, generative AI has become a conversational or interactive tool as well as a task-based one, with almost 40 per cent of users reporting conversational use. When asked to characterise their relationship with generative AI, nearly half describe it as an “information resource,” with smaller groups likening it to a personal assistant or collaborator. Awareness of embedded AI is markedly higher among those who use generative AI, suggesting that experience shapes recognition of its presence in digital environments.

These patterns indicate not only rapid uptake but also emerging differences in how Australians integrate generative AI into work, study, and home life—differences that may influence both capability development and perceptions of the technology as it continues to evolve.

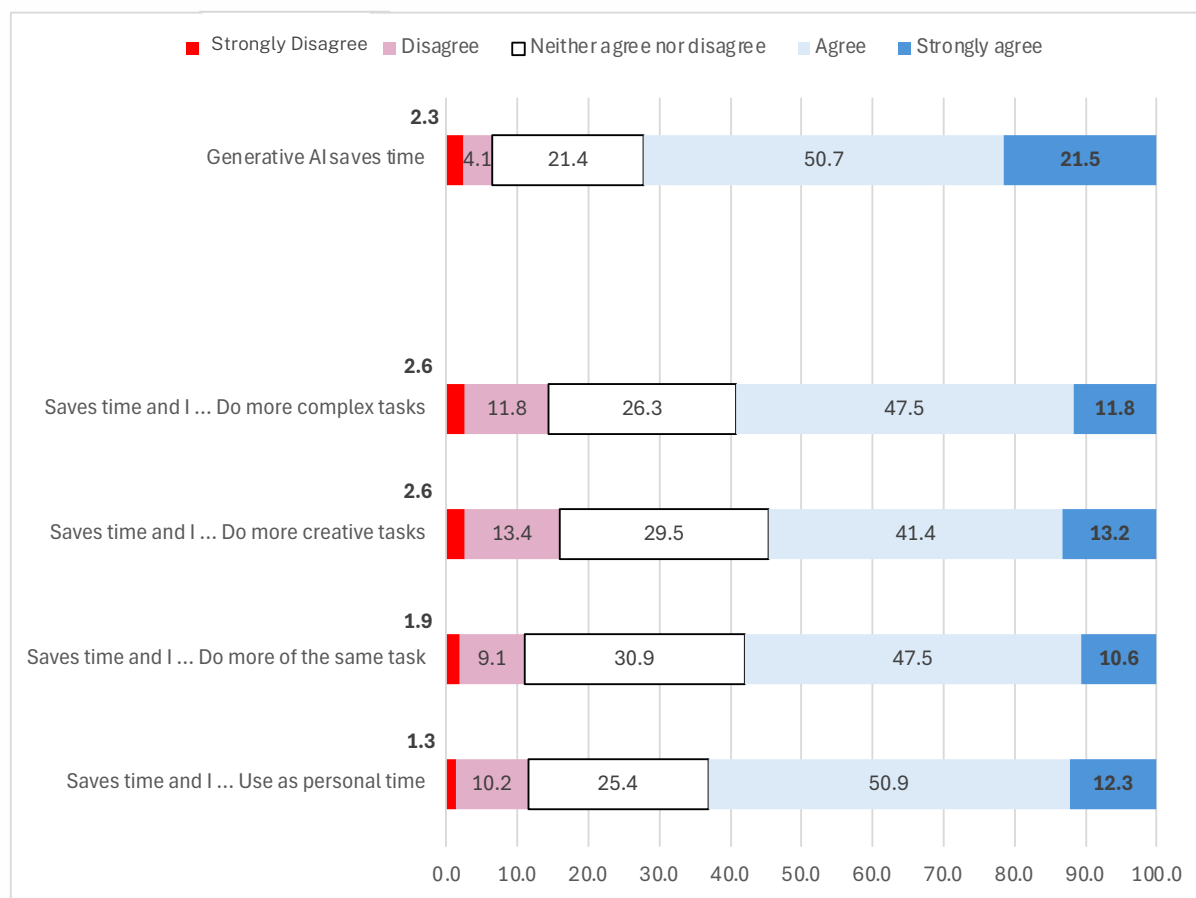
2.4 Perceived impact of Generative AI

Individual impacts

Respondents that have used Generative AI were asked about the extent to which they agree or disagree that Generative AI saves them time. Most users of Generative AI either strongly agree (21.5 per cent) or agree (50.7 per cent) that Generative AI saves them time.


Respondents that thought that Generative AI saved them time were then asked what they did with that time, using an agreement/disagreement rating across four possible options. Respondents were able to select that they did multiple things with their time. The most common response (with 63.2 per cent agreeing/strongly agreeing) was that they used the time they saved for their free personal time. There was also a high level of agreement that they did more complex tasks (59.3 per cent), did more of the same task (58.1 per cent) or did more creative tasks (54.6 per cent)

Figure 2.5 **Generative AI and saving time**



Those respondents that were employed (and had used Generative AI) were asked ‘To what extent do you agree or disagree that Generative AI improves how well you do your job?’ A slight majority of respondents either agree (39.9 per cent) or strongly agree (11.9 per cent) that Generative AI improves how well they do their job (51.9 per cent in total). We also asked students the extent to which they thought Generative AI helped them as a student. A somewhat higher per cent agreed (43.4 per cent) or strongly agreed (16.4 per cent) that Generative AI helped them as a student (59.9 per cent in total).

Figure 2.6 Perceived impact of Generative AI as an employee or student




There is substantial variation in the perceived impact of generative AI across occupation types. Respondents were asked to identify their occupation from one of eight classifications (based on the Australian and New Zealand Standard Classification of Occupations, or ANZSCO). A large number of respondents (n=124) gave a response of 'other', and were then asked to specify through text input. These responses are not included in this part of the analysis.

Remembering that the question on the survey on the perceived impact of Generative AI is only asked of those that have used it, Figure 2.7 gives the per cent of workers that agreed or strongly agreed that Generative AI improves how well the person does their job, separately for the eight stated occupations. The figure also gives the per cent of people in that occupation that have used generative AI.

There are quite large standard errors for many of the occupations, due to small sample sizes. However, the key difference appears to be between managers and professionals on the one hand, and other occupations on the other. The differences between these two broader groupings are large and statistically significant. 70.3 per cent of managers and professionals said that they had used Generative AI, compared to 51.8 per cent of other workers. Of those that have used Generative AI, 61.0 per cent of Managers and Professionals agreed or strongly agreed that it improved how well they did their job, compared to 43.1 per cent of other workers.

Figure 2.7 Perceived impact of Generative AI as an employee, by occupation



Note: The “whiskers” on the bars indicate the 95 per cent confidence intervals for the estimate.

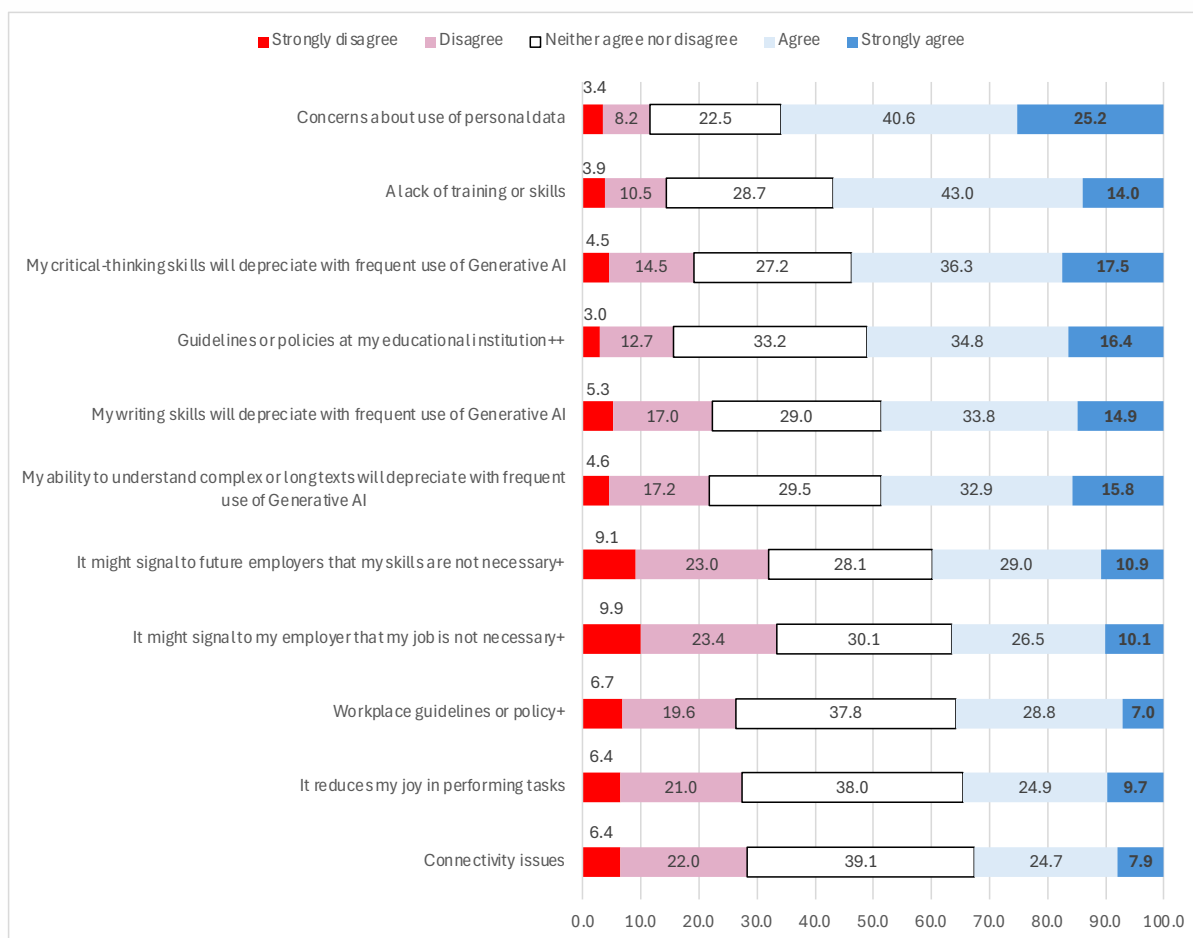
There are a number of factors that might limit a person’s use of Generative AI. Some of these factors relate to regulatory or procedural barriers, others to individual capacity constraints, and a third set of barriers relate to perceptions of negative impacts or limitations of the current systems. We asked respondents, ‘Even if you do not currently use Generative AI, to what extent do you agree or disagree that each of the following limits your use of Generative AI?’, with eleven possible limiting factors.

Figure 2.8 ranks the limitations given by respondents, based on the per cent of Australians that either agree or strongly agree. In the figure, three of the factors (labelled with a ‘+’) are asked only of those currently employed, with an additional factor (labelled with a ‘++’) asked only of current students.

The biggest limit expressed by Australians in terms of use of Generative AI is concerns about the use of personal data. Almost two-thirds of Australians agree this places a limit on their use (40.6 per cent agree, 25.2 per cent strongly agree). There were also more than half of Australians that agreed/strongly agreed that they were limited because of a lack of training or skills (57.0 per cent) or because they felt their critical thinking skills will depreciate with frequent use (53.8 per cent)

For all the statements, there were more people that agreed than disagreed. However, there were relatively few people that felt that connectivity issues were a limit (32.6 per cent agreed, 28.3 per cent disagreed) or that the limit was that it would reduce joy (34.6 per cent agreed, 27.4 per cent disagreed). The three questions that were only asked of those employed also had a relatively low agreement rate

Figure 2.8 Limits on use of Generative AI



Note: Questions that were asked only of those employed are marked with a +, those asked only of those who were studying are marked with a ++


We analysed the factors associated with these limits on Generative AI usage, focusing on the first three options that had the strongest agreement rate. Females and older Australians had the highest agreement rates regarding concerns about personal data, and a lack of training or skills. There were no differences by education, location, or country of birth, but those that were employed had a lower agreement rate for the first two limitations than those that weren't employed

There were fewer factors that were associated with agreement that depreciation of critical thinking skills was an issue limiting Generative AI use. The only variable that was statistically significant in the model was gender, with females being more likely to agree/strongly agree than males (56.1 compared to 51.4 per cent).

Societal impacts

Respondents were also asked about their attitudes towards Generative AI adoption in broader society, and in particular the role of Generative AI skills. In general, Australians were much more likely to agree than disagree that skills in Generative AI will be an important part of work and studies. There was particularly strong agreement that Generative AI will be an important part of tertiary studies (42.7 per cent agreement and 8.0 per cent strongly agreeing).

Figure 2.9 Perceptions of importance of skills in Generative AI



The socioeconomic/demographic factor that is most predictive of views towards the importance of skills in Generative AI is education. Across all three of the domains (getting a good job, school curriculum, and tertiary curriculum) and with or without controlling for other characteristics, those with higher levels of education are more likely to agree/strongly agree that skills in Generative AI will be important.

For example, 34.1 per cent of those that have not completed Year 12 or a post-school qualification think that skills in Generative AI will be important for getting a good job in the future. This increases to 42.6 per cent of those that have completed Year 12 or have a non-degree qualification, and to 54.5 per cent for those that have a degree.

Summary and implications

Overall, Australians who use Generative AI report clear and tangible individual-level benefits, particularly in the form of time savings and improvements to work and study performance. Most users say that Generative AI saves them time, and among those who do, the majority report using that time for personal activities or to take on more complex, creative, or additional tasks.

More than half of employed users believe that Generative AI improves how well they do their job, with an even higher share of students reporting positive effects on their studies. Perceived benefits, however, vary by occupation, with managers and professionals both more likely to use Generative AI and more likely to report that it improves their job performance.

At the same time, many Australians—including those who use the technology—identify significant limits on their use, led by concerns about personal data, lack of training or skills, and fears that frequent use could erode critical-thinking ability. These findings suggest that while the direct benefits of Generative AI are already being realised by many Australians, capability gaps, trust concerns, and uneven occupational patterns may constrain broader and more equitable

uptake, highlighting the need for targeted training, clearer guidance, and support for confident, informed use across the workforce and education sectors.


2.5 Generative AI and governance

One of the more challenging debates regarding Generative AI is who should be responsible for governing its use. We asked respondents to what extent they agreed or disagreed that different actors in the Generative AI ecosystem should be responsible. Figure 2.10 shows that Australians think the responsibility should be shared across the ecosystem, but that there is a particularly important role for government and the individual.

Around two thirds of Australians either agree (40.0 per cent) or strongly agree (26.2 per cent) that ‘It is the Government’s responsibility is to govern the use of Generative AI.’ More than half of respondents (61.0 and 59.1 per cent) also agree that ‘Individual organisations should be responsible for governing their use of Generative AI’ or that ‘Individual users should be responsible for governing their own use of Generative AI.’

Although the lowest level of agreement is for an independent authority to govern the use of Generative AI (48.1 per cent agree or strongly agree), there are also a large number of people who neither agree nor disagree with the statement that ‘It should be an independent authority that is responsible for governing the use of Generative AI.’ The two options with the greatest level of disagreement is for technology developers (23.2 per cent disagree or strongly disagree) or technology providers (23.1 per cent).

Figure 2.10 Who should be responsible for governing the use of Generative AI




Potentially related to the support for government taking responsibility for policy related to Generative AI, Australians have a high level of concern for the use of Generative AI in politics.

Specifically, more than a third (35.1 per cent) said that they were very concerned, with a further 44.3 per cent saying they are somewhat concerned about Generative AI in politics (Figure 2.11).

Figure 2.11 also shows that there are some small differences in the level of concern regarding Generative AI in politics, with males being slightly less likely to be very/somewhat concerned (78.1 per cent) than females (80.7 per cent). There are much greater differences in concern by age. Younger Australians are far less likely to be concerned (72.8 per cent very/somewhat concerned for those aged 18 to 34) than those aged 35 to 54 (77.3 per cent), who are in turn less likely to be concerned than those aged 55 years and over (86.8 per cent)

Figure 2.11 Level of concern about Generative AI in politics, by gender and age




The differences by gender and age in concern about Generative AI in politics hold when we analysed using an ordered probit model. Using the same dependent variable, when we add whether or not a person has used Generative AI, we find a negative and statistically significant relationship. Specifically, 74.6 per cent of those that have used Generative AI are somewhat/very concerned, compared to 83.9 per cent of those that have not used it.

We asked respondents about three specific aspects of Generative AI during election campaigns – two potentially negative and one potentially positive. Australians were most likely to agree that Generative AI poses a risk to the democratic process by spreading misinformation (32.4 per cent strongly agree, 41.2 per cent agree). Australians also agree that Generative AI creates a risk that sensitive political data is leaked to the wrong hands (25.3 per cent strongly agree, 41.9 per cent agree).

There is a net positive agreement for the view that Generative AI provides less resourced candidates with a more equal playing field to deliver effective campaigns. However, there is much less agreement (13.5 per cent strongly agree, 31.0 per cent agree) than there is for the more negative statements related to risk.

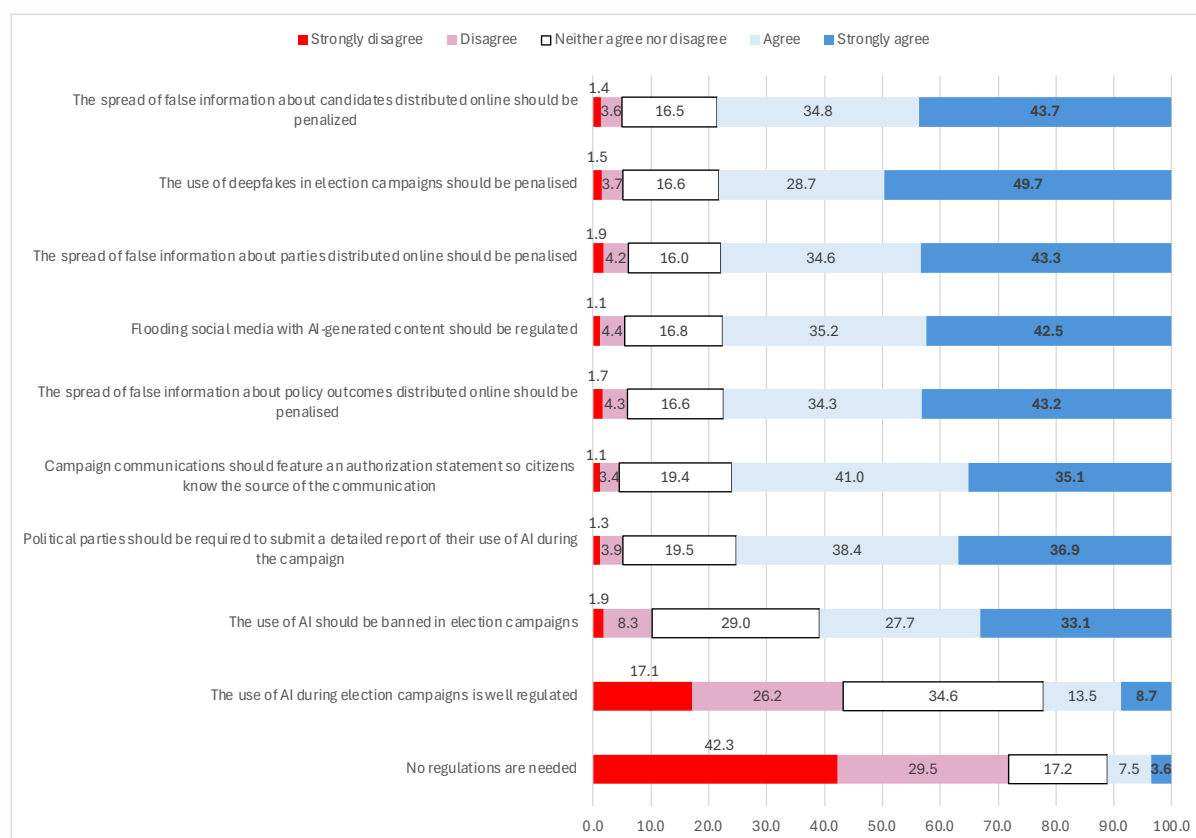
Figure 2.12 Level of concern about specific aspect so Generative AI in politics



We asked respondents about specific AI regulations in election campaigns. Two of the options had very low acceptance rates with only 11.1 per cent of Australians agreeing or strongly agreeing that no regulations are needed, and 22.1 per cent agreeing or strongly agreeing that the use of AI during election campaigns is well regulated.

There was general agreement (60.8 per cent agreeing or strongly agreeing) that the use of AI should be banned in election campaigns. However, there was even greater agreement for the seven specific and somewhat more nuanced policy responses, with at least three-quarters of Australians agreeing or strongly agreeing with each of the proposals.

Figure 2.13 Support for AI regulations in election campaigns




The policy view that has the most even distribution across the given possible response options is whether or not ‘The use of AI during election campaigns is well regulated.’ We analysed the factors associated with disagreement/agreement with this statement, with the largest differences being by age, and language spoken at home. 29.1 per cent of those aged 18 to 34 agreed or strongly agreed that the use of AI during election campaigns is well regulated. Only 23.1 per cent of those aged 35 to 54 years agreed/strongly agreed, and an even smaller percentage of those aged 55 years over agreed/strongly agreed (15.4 per cent). Those that speak a language other than English at home are far more likely to agree/strongly agree that AI during elections is well regulated (30.7 per cent) than those that speak English only (20.4 per cent).

The final attitudinal question we asked about Generative AI related to trust in institutions. We asked about trust in eight institutions ‘to develop or manage Generative AI responsibly, with response options ranging from trusting those institutions completely to not trusting them at all. The most trusted institution that we asked about was universities and research institutions. More than two-thirds of Australians either trusted universities/research institutions completely (9.2 per cent) or somewhat (59.2 per cent)

The least trusted institutions were the American government (24.8 per cent trusted at least somewhat) and particularly the Chinese government (17.4 per cent). Tech companies have a slightly higher level of trust than these two foreign governments, but there is nonetheless less than half of Australians (43.3 per cent) that trust tech companies either somewhat or completely to develop or manage Generative AI responsibly

Figure 2.14 Trust in institutions to develop or manage Generative AI responsibly



The second most trusted institution that we asked about was the Australian government, with 55.8 per cent of Australians either trusting completely (7.5 per cent) or somewhat (48.3 per cent). We analysed the factors associated with trust in the Australian government. Younger Australians had a greater level of trust than those in the middle part of the age distribution, or older Australians. The clearest distinction is between those aged 18 to 34 years (62.3 per cent trusting somewhat or completely) and those aged 35 years and over (53.0 per cent).

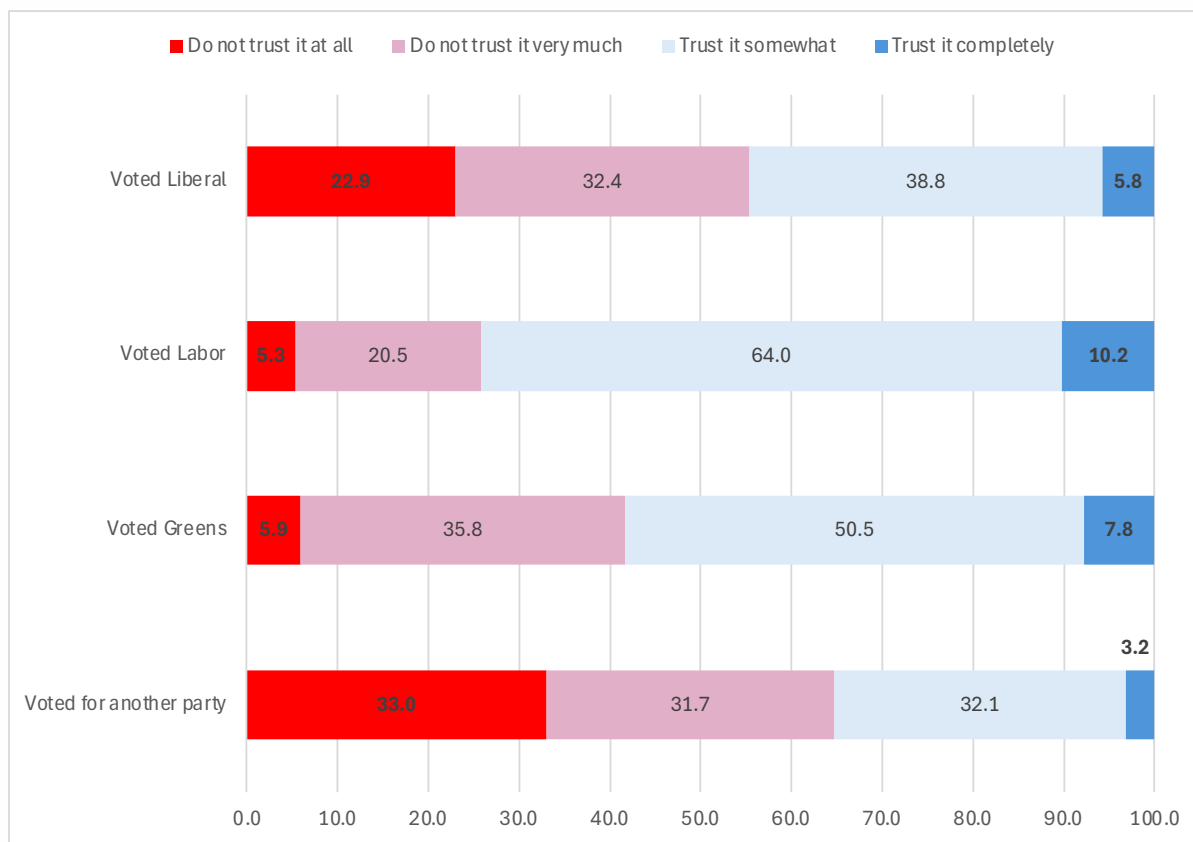
There are also significant differences by education. The difference between those that have not completed Year 12 or a post-school qualification and those that have is reasonably small (48.8 per cent compared to 53.1 per cent). However, those with a degree have a much higher level of trust than the rest of the population (65.0 per cent).

There is also a difference in trust by location. Amongst those that live in a capital city, 60.0 per cent trust the Australian government somewhat or completely to develop or manage Generative AI responsibly. This is much greater than those that live outside of a capital city (48.5 per cent).

These socioeconomic/demographic differences mirror, but are not explained by voting patterns. That is, when we include the party that a person voted for in the last election as an additional explanatory variable, the basic patterns outlined above remain. However, there are large partisan differences that are also observed in the model and the raw data.

Those that voted for the Labor party at the last election have the greatest level of trust in the Australian Government, with 74.2 per cent trusting somewhat or completely. Greens voters also have relatively high levels of trust (58.3 per cent). The lowest level of trust is for those that voted for someone else other than the major parties (35.3 per cent), with Liberal voters also reporting lower levels of trust (44.6 per cent trusting somewhat or completely).

Figure 2.15 Trust in the Australian government to develop or manage Generative AI responsibly, by party voted for in 2025 Federal election



Summary and implications

Australians see the governance of Generative AI as a shared responsibility, though with a particularly prominent role for government. Around two-thirds of respondents agree or strongly agree that governing Generative AI is the responsibility of government, and a majority also believe that individual organisations and individual users should take responsibility for their own use.

Support is lower for an independent authority, with fewer than half agreeing and many neither agreeing nor disagreeing, while technology developers and providers attract the highest levels of disagreement. These views appear linked to broader concern about the political use of Generative AI, with more than one-third of Australians very concerned and a further 44 per cent somewhat concerned about its role in politics.

Together, these findings indicate strong public support for active regulatory oversight, but also recognition that effective governance requires coordination across multiple actors in the Generative AI ecosystem. The high levels of concern about political applications underscore the importance of timely, transparent, and credible governance frameworks to sustain public trust as Generative AI becomes further embedded in Australian society.

2.6 Experiments in acceptable use of Generative AI

To test for the acceptability of using Generative AI in a professional context, we set up a controlled survey experiment with a hypothetical scenario that we varied a set of characteristics across. Specifically, we began with the following statement:

'In recent years, **Generative AI is increasingly used to generate documents and text**. We now present a **hypothetical situation involving Generative AI**. Please read the text carefully and afterwards we would like to ask you some questions.'

The vignette was structured as follows:

Imagine that <insert name from spreadsheet> created a document that you read. <She/he> is a <PROFESSION>, is <AGE>, and <Familiarity>. Now imagine that <PROPORTION> of this document is generated by AI.

The names were randomly assigned to signal gender and race/ethnicity, based on the names used in Booth et al. (2012) in their analysis of discrimination in the labour market.⁸ There were 20 female names and 20 male names. Of each of these, five names signalled an Anglo-Celtic identity, and the other 15 signalled a minority identity. The distribution of values across the other four randomly assigned variables are as follows:

- <PROFESSION>:
 - Journalist - 20%
 - Academic Researcher - 20%
 - Government Official - 20%
 - Medical Doctor - 20%
 - Lawyer - 20%
- <PROPORTION>:
 - most - 25%
 - about half - 25%
 - some - 25%
 - a small proportion - 25%
- <FAMILIARITY>:
 - you have known this person for many years - 50%
 - you have not met this person before - 50%
- <AGE>:
 - in their 20s - 33%
 - in their 30s or 40s - 33%
 - in their 50s or 60s - 33%


Although the real interest in the analysis is in the differences in values by the randomly assigned characteristics (and characteristics of respondents), the levels also have some interest (Figure 2.16). Respondents are evenly split on whether the scenario represents an acceptable use of Generative AI. However, slightly more Australians strongly disagree/disagree than strongly agree/agree that it is an acceptable use (35.1 compared to 29.3 per cent respectively).

More than three-quarters of Australians either strongly agree (31.6 per cent) or agree (45.7 per cent) that the document in the vignette should be labelled as being produced by Generative AI. There are also substantially more Australians than either strongly agree/agree that the person in the vignette is not doing their job properly (41.5 per cent) than strongly disagree/disagree (19.3

⁸ Booth, A.L., Leigh, A. and Varganova, E., 2012. Does ethnic discrimination vary across minority groups? Evidence from a field experiment. *Oxford bulletin of economics and statistics*, 74(4), pp.547-573.

per cent). In our last question, there were far more Australians that strongly agree/agree that this type of use should be regulated by government (61.4 per cent) than strongly disagree/disagree

Figure 2.16 Views on use of Generative AI in a hypothetical scenario



To analyse the impact of vignette and respondent characteristics, we estimate four ordered probit models (Appendix Table A.10). In the model are variables that capture the randomly assigned aspects of the vignette, as well as the characteristics of the respondents. As this is a hypothetical vignette with random assignment, we can be very confident in the causal direction of the vignette characteristics. However, we can also be reasonably confident in the causal direction of the observed variables as well, as it is highly unlikely that the views on the vignette impact on a person's age, sex, education, etc.

Across all four questions, we find no association between the randomly assigned gender or ethnicity of the actor in the vignette, at least as signalled by the names given. The age of the vignette actor also did not have a significant effect on any of the answers (at the 5 per cent level of significance)

There was no evidence that the occupation in the vignette impacted on whether it was seen as an acceptable use of Generative AI, or whether it should be labelled. However, compared to the base case of a journalist, respondents were less likely to agree that a government official or a medical doctor was not doing their job properly, but more likely to agree that for these occupations this use of Generative AI should be regulated by government.

The proportion of the document that is generated by AI had an impact on whether it was viewed as an acceptable use of Generative AI (smaller proportion, more acceptable), whether it should be labelled (smaller proportion, more agreement it should be labelled), and whether the person is seen as not doing their job (smaller proportion less likely to be seen as no doing their job). However, there was no statistically significant impact on whether the respondent thought that type of use of Generative AI should be regulated based on the amount of content generated.

Hypothetical familiarity with the actor had a large association. If the respondent had not met the actor before, then they were less likely to agree that the use of Generative AI was acceptable,

slightly more likely to think that the document should be labelled, more likely to think that the person is not doing their job, but no more or less likely to think the government should regulate that type of Generative AI usage.

We also considered the extent to which a respondent's own characteristics were associated with views on the acceptable use of Generative AI. As the vignette characteristics are randomly assigned, there should be no correlation between the vignette and the respondent characteristics.

Females and older Australians were less likely to think that their particular vignette is an acceptable use of AI, more likely to think the document should be labelled, and far more likely to think that type of usage should be regulated by government. The major difference by education was with regards to government regulation. Those with a degree were less likely to think that the type of AI usage in their vignette should be regulated by government. There were no significant differences (at the 5 per cent level of significance) by location, country of birth, language spoken at home, or employment status in terms of views on the vignette.

Summary and implications

In a unique survey experiment, we find that Australians are divided on whether a hypothetical vignette represents an acceptable use of Generative AI, with slightly more respondents disagreeing than agreeing that the use is acceptable. Support is far stronger for transparency and regulatory oversight: more than three-quarters believe the hypothetical document should be labelled as produced by Generative AI and more than three-fifths agreeing that this type of use should be regulated by government. Many also think the person in the vignette is not doing their job properly (41.5 per cent agree or strongly agree).

Experimental results show that several vignette characteristics shape these views: smaller proportions of AI-generated content increase perceived acceptability, reduce perceptions that the person is shirking their job, and reduce support for labelling. Familiarity also matters—respondents who had “not met” the actor in the vignette were less likely to view the use as acceptable and more likely to want the content labelled or to think the actor was not doing their job. Respondent characteristics also play a role: women and older Australians were less accepting of the vignette use and more supportive of labelling and government regulation, while those with a degree were less likely to support regulation.

Together, these findings suggest that public attitudes toward “acceptable use” are sensitive to both contextual cues and individual characteristics, with strong underlying preferences for transparency and formal oversight. The consistently high levels of support for labelling and regulation indicate that Australians expect clear disclosure when Generative AI contributes to professional outputs and want safeguards to ensure appropriate use. As organisations and government agencies expand their use of Generative AI, these expectations will likely shape both policy design and public trust, highlighting the importance of communication, clarity of standards, and responsible deployment practices.

3 Interviews

This section of the report compliments the previous survey data by drawing on in-depth interviews to provide the qualitative texture needed to understand not just **how** Australian organisations are adopting AI but **why** they are doing it.

Here, we present an analysis of AI adoption across individuals in various organisations and sectors, exploring how these levels might interact and influence one another, as well as emergent themes. We used semi-structured interviews to elicit expert views on the trends, challenges, and opportunities of AI adoption for their sector in Australia. Our study interviewed senior level individuals (n=15; f=6, m=9) from a range of organisations across Industry & Research, Education, and Government. We recruited leaders working in industries we have identified as early adopters of generative AI technologies. Interviewees were recruited through our research team's existing networks. Interviews were recorded and transcribed. Each interview took approximately 1 hour in duration (see interview questions in Appendix 2). Each participant was interviewed only once.

Overall, nine organisations were represented including: the Department of Infrastructure, Transport, Regional Development, Communications, Sport and the Arts, the Department of Employment and Workplace Relations (DEWR), consulting company KPMG, Australia's national science agency (CSIRO), the National Film and Sound Archives (NFSA), the Australian National University (ANU), and several private technology companies.

To undertake this analysis, we collected interview data that was recorded and transcribed using Zoom. Transcripts were manually cleaned and de-identified in Microsoft Word to ensure data quality and confidentiality. Cleaned transcripts were then imported into NVivo 15: qualitative data analysis software. NVivo was used to code relevant sections of the interviews according to latent themes that emerged from the interview data. Basic comparisons were done in Nvivo and then visualised in Excel. Further visualisation and deeper analysis was done in Matlab R2025b using the Statistics Toolbox.

Our interview analysis captures multiple dimensions of AI adoption, including: perceived capabilities, drivers, barriers, critical discourse. This enables a view of how and why AI is being adopted in Australian organisations across Industry & Research, Education, and Government, as well as some of the latent factors that may shape adoption behaviour.

3.1 Tool Use and Practices

Across our interviews, individuals reported using a diverse set of AI tools in their everyday workflows. These included: **AutoGen AI, Azure, ChatGPT, Claude, Co-Pilot, Databricks, Gemini, Grammarly, Meta AI (chat), Meta Llama models, Otter AI, and retrieval-augmented generation (RAG) models**. This spread reflects both general-purpose AI applications and more specialised models used for coding, data handling, or document generation.

Table 3.1 illustrates how tool use varies by sector. In our interviews, individuals working in Industry & Research reported the widest range of AI tools, spanning both mainstream generative models and more technical platforms such as Databricks and RAG models. In contrast, Education and Government participants referenced a smaller set of tools, with usage concentrated around ChatGPT, Claude, and Co-Pilot. This suggests that Industry & Research settings may be more conducive to experimentation with multiple AI systems, or that their tasks may require a broader set of capabilities.

Table 3.1: Breakdown of commonly used tools by sector

	Type Count	Commonly used AI tools
Industry, Research	9	AutoGen AI, Azure, ChatGPT, Claude, Co-Pilot, Gemini, Grammarly, Meta Llama models, and retrieval-augmented generation (RAG) models
Government	8	Azure, ChatGPT, Claude, Co-Pilot, Databricks, Gemini, Meta AI (chat), Otter AI
Education	3	ChatGPT, Claude, Co-Pilot

Note: In this table, the header 'Type Count' refers to the number of AI tool types used for a specific sector across our interviews. As Education wasn't as widely represented, this could have an effect on the number of AI tool types represented in that category. Nevertheless, the 'Industry' and 'Government' categories were quite comparable.

Beyond which AI tools are used, our interview data also highlighted how these tools are being integrated into day-to-day practices.

Figure 3.1 provides an overview through a word cloud, showing common usage types such as administration, ideation, drafting, coding, personal tasks, and proposal or report preparation. These patterns echo broader trends in AI adoption, where tools support both routine administrative tasks and higher-value creative or analytical work.



Figure 3.1. Word cloud illustrating how AI tools are used by our interviewees. Larger text indicates a stronger weighting for that usage.

A more detailed breakdown of tool-use practices across sectors is presented in Figure 3.2. This figure suggests that the Industry & Research participants we interviewed not only use a wider range of tools but also engage in a broader variety of usage types, from technical tasks such as

prototyping to strategic work, to document drafting, and brainstorming. Education and Government participants show more focused patterns of use, with fewer reported categories overall.

Together, these findings indicate that the mix of AI tools and their applications may be shaped somewhat by sector context, reflecting different workflows, expectations, constraints, and perhaps even levels of digital maturity. Industry & Research settings appear to be the most diverse and experimental in their practices, while Education and Government show more targeted forms of use aligned with their immediate operational needs. An alternative interpretation is that there may be more restrictions on what tools can be used in specific settings due to security or privacy concerns.




Figure 3.2. Breakdown of how AI tools are reportedly used by sector; across Research & Industry, Education, and Government across our interviews. On the y-axis, ‘count’ refers to the number of times that a particular usage was mentioned across all interviews. As Education wasn’t as widely represented, this could have an effect on the number of uses represented in that category. Nevertheless, the ‘Industry’ and ‘Government’ categories were quite comparable.

3.2 Motivations and Drivers of Adoption

Participants described a range of motivations for adopting AI tools, reflecting both practical needs and broader strategic considerations in their work. Many saw AI as offering capabilities that meaningfully extend what human cognition can achieve. As one individual explained:

“Obviously the human brain can't consume the kind of data points that machine learning models can. So being able to detect patterns that our brains can't actually hold, and looking at relationality in a different way in data is interesting.”

This highlights a core motivation rooted in augmentation: AI is perceived as enabling new forms of insight, pattern recognition, and relational thinking that would otherwise be inaccessible.

Others framed AI adoption through the lens of organisational value creation and economic necessity. One individual emphasised:

“The other thing is it [AI] is seen as a mechanism to drive value, to create efficiency, and if implemented well, create productivity or new value. That's a necessity for most organizations, in any industry, frankly, in this economic climate we're in.”

Here, AI is positioned not just as a tool but as an essential lever for maintaining competitiveness and delivering value under increasing economic pressure. The implication is that AI adoption is intertwined with broader organisational strategy and sustainability.

Some individuals also recognised external policy and political drivers shaping their adoption decisions. For example, one noted:

“...the government has a very clear AI acceleration agenda that's starting to emerge, which is a departure from the previous Albanese government. It's very much sold on efficiency and productivity and time and dollars saved.”

This suggests that public-sector agendas and narratives about efficiency contribute to creating momentum around AI use, particularly in government and adjacent sectors. In this context, adoption is influenced not only by internal organisational needs but also by shifts in national policy and public discourse.

These perspectives informed further investigation; how tools were used as described in interview data was coded according to latent thematic analysis and then quantified. As shown in Figure 3.3, productivity was the most common driver, accounting for 27% of all responses. This reflects the strong appeal of AI tools for streamlining tasks, speeding up drafting and analysis, and reducing administrative load.

The second most common driver was strategy (23%), suggesting that participants are also adopting AI to align with organisational goals, support innovation agendas, or maintain a competitive edge. Value (20%) and quality (14%) followed closely, pointing to the perceived benefits of AI in improving outputs, enhancing clarity, or generating more refined work.




Figure 3.3. AI adoption drivers breakdown. Productivity is shown as the largest driver of adoption proportionally (27%), followed by strategy (23%), value (20%), quality (14%), automation (10%), compliance (3%) and personal use (3%).

A smaller portion of motivations related to automation (10%), where tools were used to handle routine or repetitive tasks. The least reported drivers were compliance (3%) and personal use (3%). Compliance-related adoption generally stemmed from organisational requirements or policy shifts, while personal motivations tended to reflect curiosity or self-directed experimentation rather than formal workplace expectations.

Overall, these patterns show that AI adoption across those interviewed is driven primarily by efficiency and strategic value, with fewer instances of adoption driven by external requirements or personal interest alone. This suggests that AI tools are being integrated purposefully into work

practices, especially where they can provide tangible improvements in productivity and organisational outcomes.

3.3 Challenges & Inhibitors

While participants reported clear motivations for adopting AI tools, they also identified a range of challenges that limit or slow down adoption. These inhibitors reflect practical barriers but also deeper concerns about trust, reliability, ethics, and the broader organisational or societal impacts of AI.

A strong theme across interviews was risk awareness, particularly around sensitive data. As one individual put it:

“This is too risky. We don't want you to do this, because we don't want all these people down here uploading all our sensitive stuff to the internet.”

This quote captures both personal hesitation and organisational resistance, driven by concerns about data exposure, unclear data-processing pathways, and potential breaches. Such concerns remain a major inhibitor even when users are otherwise enthusiastic about the technology.

Another commonly reported challenge relates to the reliability and unpredictability of AI outputs. Participants recognised the impressive capabilities of AI but also its limitations:

“It's always going to be a technology which deeply needs humans, in that it can do remarkable things, but also make completely stupid mistakes.”

This underscores a central tension: AI can accelerate work, but only when users can trust the outputs. The need for human oversight is not merely a precaution; it becomes an essential part of the workflow, often reducing the time-saving benefits that AI promises.

Several individuals also expressed concerns about authenticity and professional integrity. One noted:

“The logic in it might not be robust enough, or, if someone were to challenge me on it, I would reflect that ...it doesn't accurately reflect my own position, but I've taken it because it's neatly packaged up and was an easy solution.”

This points to a subtle but important inhibitor: that even when AI use is convenient, this can introduce hesitations around AI adoption in individuals. This may be due to an individual skimming over important details and taking AI information as given. It may be the case that relying too heavily on AI-generated content may create vulnerabilities in roles where justification and personal authenticity matter.

Concerns extended beyond immediate work tasks toward environmental and societal impacts. One individual reflected:

“I'm conscious of the fact that every time I hit generate there is water, and electricity, and also a reinforcing feedback loop that says: these things are useful, please build more of them.”

This highlights a growing awareness of AI's environmental cost as well as its role in accelerating critical resource demand.

As shown in Figure 3.4 below, these qualitative insights align with the broader quantitative patterns. Security and privacy emerged as the most prominent inhibitors (18%), followed by concerns about authenticity (15%) and inaccuracy (13%). Together, these issues reflect uncertainty about whether AI outputs can be trusted and whether organisations can safely integrate these tools into existing processes.




Figure 3.4. AI adoption inhibitors breakdown. Across our interviews, the largest factor is concerns around security and privacy (18%), followed by authenticity (15%), inaccuracy (13%), environmental and cultural impacts (13%), human intervention requirements (8%), disappointment with outputs (10%), compliance anxiety (7%), and explainability issues (3%).

Participants also raised concerns about environmental and cultural impacts (13%), touching on sustainability, representational fairness, and long-term social consequences. Human intervention requirements (8%) and disappointment with outputs (10%) further highlight that AI often falls short of users' expectations, requiring extra work or producing content that must be rewritten or extensively edited.

Finally, compliance anxiety (7%) and explainability issues (3%) point to the importance of clarity, transparency, and governance. In this context, compliance anxiety is the persistent worry organisations or individuals feel about failing to meet regulatory or policy requirements for AI use, often intensified by uncertainty, evolving rules, or unclear expectations. Uncertainty about rules, or a lack of understanding of how AI reaches its conclusions, can reduce users' willingness to apply the technology in meaningful or high-stakes contexts.

Overall, these findings show that adoption barriers are multifaceted, involving a mix of ethical, technical, organisational, and experiential concerns. Even among participants who actively use AI tools, many remain cautious, emphasising the need for clearer guidance, better reliability, and stronger safeguards around data security, privacy, and responsible use. The result is a pattern where AI adoption is progressing, but not without hesitation, negotiation, and significant human oversight.

3.3.1 Overcoming Challenges & Inhibitors

The individuals we interviewed highlighted a range of strategies that help them navigate or overcome barriers related to AI adoption. A consistent theme was the importance of clear guidance, training, and regulatory frameworks. As one participant explained:

"...we have mandatory training, which provides the guardrails, ethical use, appropriate use."

This reflects a wider sentiment that structured learning and well-defined expectations reduce uncertainty and enable users to feel confident engaging with AI tools. It appears that a focus for organisations on investing in training may support safe and responsible use.

At the same time, participants described ongoing challenges associated with the quality and reliability of AI outputs. One individual noted:

“I will spend another hour or two just correcting what it has created.”

This frustration suggests that productivity, the primary driver of adoption as seen previously in Figure 3.4, can quickly erode when outputs require extensive revision. This implies that meaningful adoption is possible only when tools reduce workload burden in terms of time-savings. When they do not, users may disengage or restrict their use of AI to low-risk or low-effort tasks. This highlights the need for improved model accuracy, better domain adaptation, and clearer expectations about when AI-generated outputs are likely to be reliable.

A further theme emphasised the continuing role of humans-in-the-loop. One participant reflected:

“So, in the end, we humans have the decision. We are still in control, and we are the ones deciding all the rules and regulations.”

This reinforces the idea that AI is not seen as replacing decision-making but augmenting it. Human oversight remains central to accountability, particularly when outputs intersect with organisational policy, ethical responsibility, or professional standards. In general, this sentiment points to AI adoption that is hybrid and supervisory rather than fully automated.

Taken together, these insights suggest that overcoming AI adoption inhibitors may require a combination of ongoing training, clear governance, and continued emphasis on human agency. When users feel equipped, supported, and in control, they may be more willing to integrate AI into their workflows, even when the tools are imperfect. This positions human oversight not as a limitation, but perhaps as an enabling constraint for confident AI adoption.

3.4 Emergent Narratives

Looking deeper into the interview data, a set of narrative patterns emerged around how participants perceive AI and its impacts generally. While concerns were certainly present, the overall sentiment leaned more towards stories framing ‘**AI as Hero**’ rather than ‘**AI as Villain**’, although both narratives were present. These narratives reveal how our interviewees may position AI within their professional identities, aspirations, and anxieties, which might play an important role in shaping adoption behaviours.

AI as Hero

We categorised the emergent narrative of ‘AI as Hero’ as positive sentiment that perceived AI as ‘saving the day’ or providing capabilities that were not otherwise possible. In our interviews, individuals often described AI as a catalyst for innovation, efficiency, and creative possibility.

Participants frequently described AI as a heroic catalyst for innovation. For some, AI signified a moment of technological transformation for society in general:

“AI has the potential for us to move into, I think, a golden age... I genuinely believe we’re at the beginning of a golden age”

This sense of optimism was also observed among those working with AI for software development, where AI was thought to enable a beneficial increase in productivity. As one participant noted:

“I’m excited about using AI for making stuff with software. And I’m much less excited, and maybe even, blasé about using it for writing long documents.”

Similarly, AI was described in a positive light as accelerating prototyping and experimentation:

“for writing software, [AI is] really great, [it has] lowered the amount of effort required to knock out an MVP or a prototype in a way that’s very exciting for trying new stuff and seeing what works.”

In our interview data, the narrative of AI as Hero was framed AI as an empowering collaborative tool that expands creative capacity, speeds up development (particularly of software development and prototyping), and thus is perceived as lowering the threshold for innovation.

AI as Villain

Alongside this optimism, a counter-narrative also emerged; emphasising the risks and harms associated with AI. In our interviews, some individuals expressed concern about malicious use. For example, one individual noted:

"I feel like the leverage that it gives to bad actors is at least as great, and possibly greater than the leverage it gives to good actors."

More broadly, cultural anxieties appeared to be echoed in fears about societal decline, cognitive offloading, and perceived loss of human capability:

"...there's plenty of people who think it's the worst thing ever, and it's ruining society, and it's ruining our brains and turning them to mush."

The insights gained from these two emergent narratives: 'AI as Hero' and 'AI as Villain' builds on the existing literature that adoption is shaped not simply capability or access to tools. Instead it may be shaped by the deeper stories people use to make sense of technological change. These narratives mirror existing literature framing AI as both an enabler and a risk, producing a dynamic tension that influences how individuals experiment, where organisations invest, and what forms of readiness are required. The coexistence of these two narratives suggests that AI adoption may continue to be emotionally charged; driven as much by identity, values, and cultural context as by technical affordances. Importantly, this dual dynamic reinforces that effective AI adoption strategy must engage with both hope and concern: supporting people to harness AI's potential while addressing the uncertainties that shape its real-world use.

It is important to note that these narratives illustrating 'AI as hero' or 'AI as Villain' did not exist in isolation, and our interviewees told multiple stories depicting AI in different ways. This illustrates a perception of AI adoption as a tension between positive and negative factors.

3.5 The Future of AI Adoption: A Uniquely Australian Perspective

In this study, individuals articulated a clear sense that Australia has a unique opportunity to shape a distinctive approach to AI adoption grounded not simply in technological capability, but in values, identity, and strategic adaptation. Rather than trying to replicate large-scale efforts underway in other nations, participants emphasised the potential for Australia to lead AI adoption by aligning development with principles that reflect the country's nuanced social, cultural, and environmental priorities.

As one participant noted, this approach positions sustainability at its centre:

"...rather than being caught up on things like data and AI sovereignty, we could focus on values and principles grounded in sustainable practices"

The above quote also mentioned 'values' which often arose in our interviews. A recurring theme was the question of **whose** values, and whose version of 'Australia' should guide this development. Individuals mentioned that developing an "Australian AI" requires deeper consideration of national identity and representation. One interviewee in particular captured this tension directly:

"...we'd like something a little bit more reflective and comprehensive about what we mean by being Australian. Our nation can't have an Australian AI until we figure out what our identity is."

This suggests that global leadership in AI adoption may depend as much on alignment with cultural values (which are currently not adequately and inclusively characterised) as on technological capability.

Participants also reflected on what a national AI capability should look like in practice, particularly in relation to AI model development. There was strong agreement that national capability does not require building entirely new models, but instead working strategically with existing ones:

"I think it's about building that sovereign capability, but that does not mean building a new model from the ground up and cutting yourself off from existing models. It means working with, it means working with existing models, tools, and adapting them to our purpose."

This adaptive approach was seen as both realistic and necessary, focusing on fine-tuning, and improving existing systems for Australian contexts:

“After you train that model, can we tune this model? Can we improve this model? Can we use this model better to have an Australian sovereign AI, or innovation - I think that's absolutely what we need to do.”

These perspectives point to a model of AI adoption defined by grounding in Australian culture, and an individual and organisational commitment to values-based technological development. These views are echoed in broader discussions across Australian policy, ethics, and governance. Scholars of digital sovereignty argue that meaningful national capability emerges not from technological isolation, but from the ability to adapt and govern global systems in contextually grounded ways which are often overlooked.^{9 10} For example, rather than considering AI sovereignty as absolute power, we should think of sovereignty as embodied power, and sovereignty as institutional power.¹¹

Exactly what kinds of embodied power and institutional power AI sovereignty might provide individuals and organisations such as those who took part in our study would be worthy of future research. For example, specific discourse around Australian AI sovereignty, especially when it mirrors organisational aspirations or anxieties, may shape individual AI adoption motivation. In parallel, broader organisational desire for AI adoption may be better understood and unpacked as something produced through conversations about Australian AI capabilities, which is ultimately influenced by systems and actors that sit beyond organisational and national boundaries.

This enables us to build on existing and current policy literature around safe AI adoption¹² and begin to picture what exactly sustainable adoption might look like. This is an approach that positions Australia not as a follower in the global AI race, but as a strategically distinctive leader. For example, Australia's AI Ethics Principles¹³ emphasise values, inclusion, and trust as foundations for responsible deployment. This aligns with work on First Nations data sovereignty,^{14 15} which highlights the importance of Aboriginal and Torres Strait Islander self-governance of representation and cultural rights in national data and AI ecosystems. Together, these bodies of work reinforce the idea that a uniquely Australian approach to AI will depend as much on cultural stewardship and inclusivity as on technical infrastructure.

To build on this, we found that across our interviews, participants highlighted a core challenge: that Australia cannot meaningfully lead in AI adoption without first articulating which values, identities, and cultural priorities should guide this leadership. The tension around “whose values” and “what version of Australia” should shape AI development underscores the current absence of an inclusive, nationally shared understanding of what responsible and representative AI looks like in an Australian context.

In response to this challenge, interviewees pointed to a practical mechanism: building sovereign capability through adaptation, rather than isolation. Rather than investing in wholly new national models which remain unrealistic due to economic and capability constraints, participants argued

⁹ Couture, S., & Toupin, S. (2019). What does the notion of “sovereignty” mean when referring to the digital? *New Media & Society*, 21(10), 2305–2322. <https://doi.org/10.1177/1461444819865984>

¹⁰ Hummel, P., Braun, M., Tretter, M., & Dabrock, P. (2021). Data sovereignty: A review. *Big Data & Society*, 8(1), 2053951720982012. <https://doi.org/10.1177/2053951720982012>

¹¹ Braun, M., & Hummel, P. (2024). Is digital sovereignty normatively desirable? *Information, Communication & Society*, 28(10), 1721–1734. <https://doi.org/10.1080/1369118X.2024.2332624>

¹² Australian Government, *Australia establishes new institute to strengthen AI safety*, Department of Industry, Science and Resources, accessed 27 November 2025, <https://www.industry.gov.au/news/australia-establishes-new-institute-strengthen-ai-safety>

¹³ Commonwealth of Australia. (2019). *Australia's AI Ethics Principles*. Department of Industry, Science and Resources. <https://www.industry.gov.au/publications/australias-artificial-intelligence-ethics-principles/australias-ai-ethics-principles>

¹⁴ Kukutai, T., & Taylor, J. (2016). *Indigenous data sovereignty: Toward an agenda*. ANU press.

¹⁵ Carroll, S., Garba, I., Figueroa-Rodríguez, O., Holbrook, J., Lovett, R., Materechera, S., ... & Hudson, M. (2020). The CARE principles for indigenous data governance. *Data science journal*, 19.

for a strategy of fine-tuning, contextualising, and governing existing global models to meaningfully serve Australian needs. This aligns with international thinking on “sovereign interdependence,” which emphasises governance, contextualisation, and capability-building over attempting to achieve full technological independence¹⁶ ¹⁷. Such a mechanism supports environmental sustainability (by avoiding energy-intensive frontier model training) and leverages Australia’s strengths in regulation and cross-sector collaboration.

¹⁶ Nye Jr, J. S. (2020). Power and interdependence with China. *The Washington Quarterly*, **43**(1), 7-21.

¹⁷ Mueller, M. L. (2020). Against sovereignty in cyberspace. *International studies review*, **22**(4), 779-801.

4 Summary and Conclusions: What's Next

Together, the survey and interview components of this report have explored GAI and AI adoption across individuals, organisations, and sectors, revealing the complex interplay of readiness, motivators, challenges, and perceptions. Together, the survey-level patterns, alongside the qualitative findings capture the emotional, narrative, and cultural layers that shape how Australians understand and engage with AI, expressing both excitement and hesitation. These deeper narratives help explain not only whether individuals and organisations adopt AI, but how they might do so, and what drives momentum or resistance across different modes of use and context.

The core findings from the survey data analysis suggests that:

- Interest in emerging technologies is widespread, but unevenly distributed across the population. Younger people, those with higher levels of education, and Australians who speak a language other than English at home demonstrate substantially greater interest and more favourable attitudes towards GAI.
- Familiarity with and understanding of GAI is high, with nearly three-quarters of Australians reporting at least a basic or moderate understanding. However, only a small minority report strong understanding, and levels of familiarity vary sharply by age and education.
- GAI use is already common, with almost half of Australians having used it at least once. Use is concentrated among younger Australians, those with tertiary qualifications, and those living in capital cities, indicating an emerging capability divide.
- Among users, GAI is seen as delivering clear personal benefits, especially through time savings. Many report that it enables them to complete more complex tasks, handle larger volumes of work, or engage in more creative activities.
- Students and workers who use GAI perceive meaningful improvements in study quality and job performance, with managers and professionals the most likely to report positive impacts.
- Despite reported benefits, many Australians face barriers to confident use, including concerns about personal data and privacy, a lack of training or skills, and worries that using Generative AI might erode critical-thinking abilities.
- Australians see governance of GAI as a shared responsibility, but with a particularly strong role for government. Transparency is especially important, with clear majorities supporting labelling of AI-generated content and expressing concern about political uses of AI.
- Judgements about “acceptable use” of GAI are sensitive to context, with acceptability higher when AI contributes a smaller share of a document or task. Respondent characteristics—including gender, age, and education—also shape perceptions of appropriateness and the need for oversight

To recap the core findings of our interviews, we observed that:

- How AI tools were used varied according to organisational context; Industry reported the widest variety of use cases compared to Education and Government sectors
- The strongest motivators of adoption were perceptions that AI could enhance productivity, support strategic goals, create value, improve quality, enable automation, assist with compliance, and serve personal needs.
- The main challenges to AI adoption included concerns around security and privacy, authenticity, inaccuracy, environmental and cultural impacts, human intervention requirements, disappointment with outputs, compliance anxiety, and explainability

issues. However, many interviewees reported various strategies for overcoming these challenges.

- Further analysis revealed that AI adoption may be shaped by narratives as well as by perceived capability benefits. We observed the coexistence of dual “AI as Hero” and “AI as Villain” narratives which suggests that AI adoption is approached through a tension of optimism and concern.
- Individuals articulated a uniquely Australian pathway for AI adoption, one grounded in values, identity, and cultural context. They questioned ‘whose values’ will shape the future of Australian AI, and how it might be possible to build a type of sovereign capability through adaptation, fine-tuning, and responsible governance of existing models, reflecting an approach of sovereign interdependence rather than independence.

By integrating the survey and interview findings, several similarities emerge revealing a broadly consistent picture of an Australian public and workforce that is increasingly engaged with AI.

For instance, at a population level the survey data shows strong interest and growing familiarity, especially among younger, more educated, and urban Australians, who also report the greatest personal and professional benefits. Likewise, the Interview participants echoed these capability-oriented motivations, describing AI as a tool that enhances productivity, strategic alignment, and creative or analytical capacity. While the survey quantifies enthusiasm and perceived value, the interviews suggest that these perceptions are shaped by organisational context, professional demands, and the emotional and cultural narratives people bring to their work. That is, in both datasets, productivity emerges as a key driver, but the interviews illuminate the conditional nature of this benefit: AI improves work only when outputs are accurate, trustworthy, and do not require extensive correction.

Both survey and interview data also align in identifying barriers, though emphasis differs in character and scale. For example, survey participants flag privacy concerns, lack of training, and fears of diminishing critical-thinking skills, while interviewees describe inhibitors such as authenticity, model inaccuracies, and anxieties around compliance. The interviews suggest that individuals navigate these barriers through their own organisations’ guardrails, and quality checks, suggesting that confidence in AI is as much a product of governance environments as of individual skill or understanding. Importantly, interviews introduce narrative framings (“AI as Hero” vs. “AI as Villain”), questions of identity and values in shaping AI adoption, and reflections on what a uniquely Australian approach to AI might entail.

For organisations, these insights offer guidance on navigating AI adoption. For instance, the findings outlined here suggest that successful AI adoption may best be focused on practices that tap into key motivators such as improving productivity, strategy, and value, and address key inhibitors such as security and privacy, as well as authenticity. Organisations also have a role to play in shaping the AI narrative by fostering dialogue about Australia’s unique approach to AI Adoption and its potential; particularly around understanding the nuances of the Australian context.

At the same time, organisations could avoid common misconceptions. For example, this report has reinforced that AI adoption is not only a technical challenge, and that capability is more than just tool use, but has also demonstrated that overlooking the emotional and narrative factors that influence engagement would ignore the dynamics at play influencing AI adoption.

By recognising AI adoption as a socio-technical process, one shaped by people as much as technology, individuals and organisations can shape the conditions necessary for meaningful and sustainable integration of AI into their work.

5 References

- ASIC, 2024. *Beware the gap: Governance arrangements in the face of AI innovation*. [online] Available at: <https://asic.gov.au/regulatory-resources/find-a-document/reports/rep-798-beware-the-gap-governance-arrangements-in-the-face-of-ai-innovation/> [Accessed 14 Jan. 2025].
- Australian Government, *Australia establishes new institute to strengthen AI safety*, Department of Industry, Science and Resources, accessed 27 November 2025, <https://www.industry.gov.au/news/australia-establishes-new-institute-strengthen-ai-safety>
- Baregheh, A., Rowley, J. and Sambrook, S., 2009. Towards a multidisciplinary definition of innovation. *Management Decision*, 47(8), pp.1323–1339.
- Bell, G., 2006. The age of the thumb: A cultural reading of mobile technologies from Asia. *Knowledge, Technology & Policy*, 19(2), pp.41–57. [online] Available at: <https://doi.org/10.1007/s12130-006-1023-5>.
- Biddle, N. and ANU Centre for Social Research & Methods, 2023. *Views of Australians towards science and AI*. [online] Available at: <https://csrcm.cass.anu.edu.au/research/publications/views-australians-towards-science-and-ai> [Accessed 1 May 2025].
- Bick, A., Blandin, A. and Deming, D.J., 2024. *The rapid adoption of generative AI*. National Bureau of Economic Research Working Paper Series. [online] Available at: <https://doi.org/10.3386/w32966> [Accessed 1 May 2025].
- Bonney, K., Breaux, C., Buffington, C., Dinlersoz, E., Foster, L.S., Goldschlag, N., Haltiwanger, J.C., Kroff, Z. and Savage, K., 2024. Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey. *National Bureau of Economic Research Working Paper Series*. [online] Available at: <https://doi.org/10.3386/w32319>.
- Braun, M., & Hummel, P. (2024). Is digital sovereignty normatively desirable? *Information, Communication & Society*, 28(10), 1721–1734. <https://doi.org/10.1080/1369118X.2024.2332624>
- Brynjolfsson, E., Li, D. and Raymond, L.R., 2023. Generative AI at work. *National Bureau of Economic Research Working Paper Series*. [online] Available at: <https://doi.org/10.3386/w31161>.
- Carroll, S., Garba, I., Figueroa-Rodríguez, O., Holbrook, J., Lovett, R., Materechera, S., ... & Hudson, M. (2020). The CARE principles for indigenous data governance. *Data science journal*, 19.
- Choung, H., David, P. and Ross, A., 2023. Trust in AI and its role in the acceptance of AI technologies. *International Journal of Human-Computer Interaction*, 39(9), pp.1727–1739. [online] Available at: <https://doi.org/10.1080/10447318.2022.2050543>.
- Couture, S., & Toupin, S. (2019). What does the notion of “sovereignty” mean when referring to the digital? *New Media & Society*, 21(10), 2305–2322. <https://doi.org/10.1177/1461444819865984>
- Cui, Z. (Kevin), Demirer, M., Jaffe, S., Musolff, L., Peng, S. and Salz, T., 2024. The effects of generative AI on high skilled work: Evidence from three field experiments with software developers. *SSRN Scholarly Paper*. [online] Available at: <https://doi.org/10.2139/ssrn.4945566>.
- Dell’Acqua, F., McFowland, E., Mollick, E.R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Kraymer, L., Candelon, F. and Lakhani, K.R., 2023. Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *SSRN Electronic Journal*. [online] Available at: <https://doi.org/10.2139/ssrn.4573321>.
- Deloitte, 2024. *State of generative AI in the enterprise 2024*. [online] Available at: <https://www2.deloitte.com/us/en/pages/consulting/articles/state-of-generative-ai-in-enterprise.html> [Accessed 14 Jan. 2025].
- Dosi, G., 1988. Sources, procedures, and microeconomic effects of innovation. *Journal of Economic Literature*, 26(3), pp.1120–1171. [online] Available at: <http://www.jstor.org/stable/2726526>.
- Eloundou, T., Manning, S., Mishkin, P. and Rock, D., 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), pp.1306–1308. [online] Available at: <https://doi.org/10.1126/science.adj0998>.
- Fagerberg, J., Mowery, D.C. and Nelson, R.R. (eds.), 2004. *The Oxford handbook of innovation*. Oxford: Oxford University Press.
- Felten, E.W., Raj, M. and Seamans, R., 2023. How will language modelers like ChatGPT affect occupations and industries? *SSRN Scholarly Paper*. [online] Available at: <https://doi.org/10.2139/ssrn.4375268>.

Fletcher, R. and Nielsen, R.K., 2024. What does the public in six countries think of generative AI in news? *Reuters Institute for the Study of Journalism*. [online] Available at: <https://doi.org/10.60625/RISJ-4ZB8-CG87>.

Gillespie, N., Lockey, S., Curtis, C., Pool, J. and Akbari, A., 2023. Trust in artificial intelligence: A global study. The University of Queensland and KPMG Australia. [online] Available at: <https://doi.org/10.14264/00d3c94>.

Heimberger, H., Horvat, D. and Schultmann, F., 2024. Exploring the factors driving AI adoption in production: A systematic literature review and future research agenda. *Information Technology and Management*. [online] Available at: <https://doi.org/10.1007/s10799-024-00436-z>.

Holmström, J., 2022. From AI to digital transformation: The AI readiness framework. *Business Horizons*, 65(3), pp.329–339. [online] Available at: <https://www.sciencedirect.com/science/article/pii/S0007681321000744> [Accessed 25 Feb. 2025].

Humlum, A. and Vestergaard, E., 2024. The adoption of Chat GPT. [online] Available at: <https://www.google.com/search?q=The+Adoption+of+Chat+GPT%2C+2024.+Andres+Humlum+and+Emilie+Vestergaard> [Accessed 15 Jan. 2025].

Kukutai, T., & Taylor, J. (2016). *Indigenous data sovereignty: Toward an agenda*. ANU press.

Kuo, E.C.Y. (ed.), 2002. *Internet in Singapore: A study on usage and impact*. Singapore: Times Academic Press.

Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N., 2025. The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers. [online] Available at: <https://www.microsoft.com/en-us/research/publication/the-impact-of-generative-ai-on-critical-thinking-self-reported-reductions-in-cognitive-effort-and-confidence-effects-from-a-survey-of-knowledge-workers/> [Accessed 20 Mar. 2025].

Liu, Y. and Wang, H., n.d. Who on Earth is using generative AI? A working paper by the World Bank. [online] Available at: <https://openknowledge.worldbank.org/entities/publication/5a876bd0-f85a-479b-ae32-cf0b7f33792f> [Accessed 14 Jan. 2025].

Maslej, N., Fattorini, L., Perrault, R., Parli, V., Reuel, A., Brynjolfsson, E., Etchemendy, J., et al., 2024. Artificial Intelligence Index Report 2024. *arXiv*. [online] Available at: <https://doi.org/10.48550/arXiv.2405.19522>.

McKinsey & Company, 2025. *The state of AI in early 2025*. [online] Available at: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai> [Accessed 23 Mar. 2025].

Merhi, M.I. and Harfouche, A., 2023. Enablers of artificial intelligence adoption and implementation in production systems. *International Journal of Production Research*, 62(15), pp.5457–5471. [online] Available at: <https://doi.org/10.1080/00207543.2023.2167014>.

Mueller, M. L. (2020). Against sovereignty in cyberspace. *International studies review*, 22(4), 779–801.

Noy, S. and Zhang, W., 2023. Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), pp.187–192. [online] Available at: <https://doi.org/10.1126/science.adh2586>.

Nye Jr, J. S. (2020). Power and interdependence with China. *The Washington Quarterly*, 43(1), 7–21.

Peng, S., Kalliamvakou, E., Cihon, P. and Demirel, M., 2023. The impact of AI on developer productivity: Evidence from GitHub Copilot. [online] Available at: <https://www.microsoft.com/en-us/research/publication/the-impact-of-ai-on-developer-productivity-evidence-from-github-copilot/> [Accessed 1 May 2025].

Riskin, J., 2016. *The restless clock: A history of the centuries-long argument over what makes living things tick*. Chicago: The University of Chicago Press.

Rogers, E.M., 2003. *The diffusion of innovations*. 5th ed. New York: The Free Press.

Selwyn, N., Cordoba, B.G., Andrejevic, M. and Campbell, L., 2020. AI for social good: Australian public attitudes toward AI and society. [online] Available at: <https://doi.org/10.26180/13159781.v1>.

Semrush. (n.d.). Retrieved November 27, 2025, from <https://www.semrush.com/>

Slack, 2024. Despite AI enthusiasm, Workforce Index reveals workers aren't yet unlocking its benefits. [online] Available at: <https://slack.com/intl/en-au/blog/news/the-workforce-index-june-2024> [Accessed 15 Jan. 2025].

Straub, E.T., 2009. Understanding technology adoption: Theory and future directions for informal learning. *Review of Educational Research*, 79(2), pp.625–649. [online] Available at: <https://doi.org/10.3102/0034654308325896>.

Taherdoost, H., 2018. A review of technology acceptance and adoption models and theories. *Procedia Manufacturing*, 22, pp.960–967. [online] Available at: <https://www.sciencedirect.com/science/article/pii/S2351978918304335> [Accessed 1 May 2025].

Teo, T.S.H., Lin, S. and Lai, K.H., 2009. Adopters and non-adopters of e-procurement in Singapore: An empirical study. *Omega*, 37, pp.972–987. [online] Available at: <http://dx.doi.org/10.1016/j.omega.2008.11.001>.

The University of Queensland, 2021. *Trust in Artificial Intelligence: A Global Study*. [online] Available at: <https://ai.uq.edu.au/project/trust-artificial-intelligence-global-study> [Accessed 1 May 2025].

Uren, V. and Edwards, J.S., 2023. Technology readiness and the organizational journey towards AI adoption: An empirical study. *International Journal of Information Management*, 68. [online] Available at: <https://www.sciencedirect.com/science/article/pii/S0268401222001220> [Accessed 25 Feb. 2025].

Yang, J., Blount, Y. and Amrollahi, A., 2024. Artificial intelligence adoption in a professional service industry: A multiple case study. *Technological Forecasting and Social Change*, 201. [online] Available at: <https://www.sciencedirect.com/science/article/pii/S0040162524000477>.

Zhang, B. and Dafoe, A., n.d. Artificial intelligence: American attitudes and trends. *Centre for the Governance of AI*. [online] Available at: <https://www.governance.ai/research-paper/artificial-intelligence-american-attitudes-and-trends> [Accessed 14 Jan. 2025].

Zhu, K., Kraemer, K.L. and Xu, S., 2006. The process of innovation assimilation by firms in different countries: A technology diffusion perspective on e-business. *Management Science*, 52, pp.1557–1576. [online] Available at: <https://doi.org/10.1287/mnsc.1050.0487>.



Australian
National
University



Capabilities

Remembers what user said
earlier in the conversation

Allows user to provide follow-
up corrections

Trained to decline inappropriate
requests

Limited knowledge