



Inquiry into the use of generative AI in the education system

Submission to the House Standing Committee on
Employment, Education and Training

14 July 2023

Executive Summary

The University of Melbourne welcomes the opportunity to contribute to the inquiry into the use of generative artificial intelligence (AI) in the Australian education system by the House Standing Committee on Employment, Education and Training.

The advent of widely accessible and user-friendly generative AI tools has led many to speculate on the ways in which it will transform our society for both good and ill. Many are fearful of the consequences of these tools, which have generated widespread discussion about the nature of work, creativity and even what it means to be human.

For the higher education sector, generative AI poses both opportunities and risks. Students will inevitably come across generative AI during their studies. Moreover, these tools are becoming increasingly ubiquitous, embedded into everyday programs and apps such as word processors and search engines. Research processes across disciplines may be modified to reflect the availability of generative AI tools. It is likely that almost all workforces and professions will be impacted by generative AI in some way. It is therefore crucial that staff and students understand how to use generative AI effectively and ethically.

Generative AI can be a helpful tool for teaching, allowing students to develop their critical thinking and evaluative judgement skills by generating outputs from platforms such as ChatGPT and then improving and critiquing that output. Already, some subjects at the University of Melbourne involve the use of generative AI in their assessments. There are also opportunities for researchers to use these tools when undertaking their research, using them to develop code, help design surveys, and test methodologies.

However, these opportunities must be balanced against the need for academic and research integrity. For example, at the University of Melbourne, the use of generative AI tools to produce work that is subsequently submitted for assessment is considered academic misconduct, unless permission to use such tools has been given by academic staff and their use has been acknowledged by the student. Similarly, generative AI can only be used in research outputs where the material generated is acknowledged.

There are also broader challenges inherent in the tools that must be considered. For example, generative AI tools such as ChatGPT are known to produce convincing but false information, a tendency that is referred to by developers as “hallucinating”. Additionally, generative AI tools are fundamentally biased, replicating and reinforcing the biases present in the human-created materials on which they are trained. This means that any AI-generated outputs require critical evaluations by humans. Moreover, Universities and other education providers will need to ensure students have equitable access to these tools, where they are required for education and research.

Universities will inevitably need to adapt their offerings to respond to generative AI. This will include reconsidering assessment to include more authentic, continuous forms of assessment where the opportunity for academic misconduct is reduced. Universities will need to ensure Higher Degree by Research (HDR) candidates, for example, understand the importance of generating their own content, noting that writing is intimately linked with cognition. Universities will also need to ensure their students and staff are AI literate.

Universities should continue to have autonomy over their own institutional policies and responses to generative AI. However, the Government may wish to develop optional guidance on the use of generative AI in the sector. The Government could also establish communities of practice to share ideas, best practice and lessons learned in generative AI in higher education.

Changes to the Higher Education Standards Framework may not be required at this stage, noting that they are currently broad enough to address academic and research integrity risks associated with AI. However, the Government should encourage transparency in the use of generative AI in core functions within Universities.

It will also be important for the Government to invest in research and workforce development in AI and related technologies, to ensure Australia can develop its own generative AI resources. Separately, the Government may wish to consider establishing a research centre to further investigate generative AI in education.

For more information or to discuss the submission, Professor Gregor Kennedy, Deputy Vice-Chancellor (Academic) can be contacted at gek@unimelb.edu.au.

Recommendations

The University of Melbourne recommends that:

1. The Australian Government develop optional guidance for the sector on the use of generative artificial intelligence (AI) in teaching, learning, assessment, and research. For example, the Government could outline the various practical and ethical issues (including data privacy, biases, accountability, and reliability of generative AI) that institutions may wish to consider when developing and updating their internal policies. However, Universities should continue to have autonomy over their own generative AI policies to ensure they are appropriate for their communities.
2. The Australian Government work with the sector to establish communities of practice on generative AI in education between Universities, Government and industry partners. This could include conferences, workshops, and forums where stakeholders exchange ideas, best practice, and lessons learned in generative AI in higher education. Participants could discuss and share alternative forms of assessment that reduce the risk of academic misconduct, creative ways to integrate generative AI into teaching and assessment, and emerging ethical issues.
3. TEQSA continue to develop guidance for the sector and share useful resources on AI. At this stage, updates to the Threshold Standards to reflect generative AI may not be necessary, as the Standards already require institutions to effectively assess learning outcomes and mitigate risks to academic and research integrity. This may change as the technology advances, however.
4. The Australian Government encourage transparency in the use of generative AI in core student- and staff-facing functions within Universities and in other education institutions.
5. The Australian Government and industry invest in R&D and workforce development in AI and related technologies, noting the risks for Australia with generative AI resources being mainly concentrated within US-based private companies.
6. The Australian Research Council consider establishing a fund or Centre of Excellence for research into generative AI in education.

Response to inquiry

Background

It is critical that discussions of the topic are precise about what is meant by generative artificial intelligence (AI). As noted in the Rapid Response Information Report commissioned by the National Science and Technology Council, generative AI can create new content, including texts, images, and code, based on a prompt provided by a user.¹ Currently, the most well-known generative AI tool is ChatGPT, a chatbot built on a large language model (LLM) that was developed by OpenAI and launched in November 2022. There are also non-text-based generative AI tools such as DALL-E 2, which produces images and art, and MusicGen, which creates music.

LLMs are trained on massive collections of text collected from the Web. Internally, LLMs are nothing more than very large collections of text fragments — word and subword combinations. In response to a user's prompt the generative algorithm assembles the best-matching fragments from its memory and starts producing an answer. With this partial sentence, it then dips into memory again to find fitting continuation fragments and outputs the most apt one as its next segment, continuously extending its sentences as the output unfolds.

There is no inner logic, no reasoning, no repository of knowledge, and no explicit guidance to anything it does except the numerical match scores that provide segment continuation rankings. It simply recognises patterns in data and produces answers based on those patterns. However, its fragment selection algorithm enforces fluency, and that alone makes it seem more “intelligent” than it is.

The input material is drawn from a very large collection of online documents: accurate, inaccurate, hypothetical, polemic. Similar fragments are collapsed and stored in the same places in the LLM, making subsequent identification of the source of any fragment impossible to determine. Two fragments that are individually true may be combined to produce a false statement without the LLM having any awareness of that fact.

The seemingly incredible abilities of ChatGPT to answer questions, produce poetry and humour, speak many languages, generate software code, summarise and review papers, and more, has become tempered with the uncomfortable realisation that its outputs often include inaccurate or misleading information, may switch sides to opposing points of view on a single command, and is limited to the kinds of material the model was trained on. The same limitations apply to Bard (its equivalent at Google) and Bing Search's Sydney (Microsoft's version). Generative AI tools also are known to produce harmful instructions or biased content (explored later in this submission).

Like all disruptions, the emergence of generative AI creates not only challenges but also creates opportunities. These must be considered thoughtfully and balanced through institutions' policies, sector-wide guidelines, and government regulation and legislation.

Strengths and benefits

Students will increasingly be exposed to generative AI during their studies. This exposure will come not only through purpose-built tools, but also through AI's proliferation into ordinary workplace tools such as Microsoft Office and Google Suite. Generative AI is likely to impact almost all disciplines, workforces, and professions. Universities should therefore respond to ensure students are equipped for this. Universities must provide general and discipline-specific opportunities for students to understand and critically interrogate the risk and opportunities of generative AI— including the impacts of bias and misinformation.

Opportunities to improve teaching, learning and assessment

Generative AI tools have impressive capabilities that can be leveraged by academics to improve various aspects of teaching, learning and assessment for students. Creative ideas in this domain are beginning to emerge both globally and at the University of Melbourne. Some ideas for enhancing teaching and learning activities and assessment tasks using generative AI are summarised below:

¹ Bell, G., Burgess, J., Thomas, J., and Sadiq, S. (2023), *Rapid Response Information Report: Generative AI - language models (LLMs) and multimodal foundation models (MFMs)*, Australian Council of Learned Academies, p.2

- Students' expectations and understanding of assessment standards could be strengthened by academics using generative AI to create exemplars of varying quality for an upcoming written assessment task, and then asking students to critique and evaluate these exemplars using a marking rubric.
- Students' critical and analytical thinking and evaluative judgement skills may be developed by intentionally using generative AI tools. For example, ChatGPT could be used to generate an outline or early response to a question posed by an academic. Students could be asked to critically evaluate this response, further develop this response, and clearly demonstrate how they have improved and altered the response over time (e.g., using tracked changes in Word).
- Generative AI may be able to support accessibility and inclusion of students with disabilities, including through creation of automated captioning/audio descriptions and accessibility testing of teaching and learning materials.

Many early-adopter academics are already incorporating the use of generative AI tools in their teaching practice. This recognises the ongoing role that generative AI is likely to play and the need to equip our students with the skills and ability to work effectively with AI. The use of new generative AI tools within our pedagogical practices is in its infancy, and as the tools evolve rapidly, so will the ways in which they are incorporated.

Already, there have been early examples of adoption and use of generative AI in the classroom at the University of Melbourne. For example, students in a graduate software engineering course are using a ChatGPT-enabled process to undertake peer code review. Peer code review is a best-practice approach used by leading engineering firms but can be confronting for students and time-consuming to implement. In creating a ChatGPT-enabled approach this overcomes these issues and builds AI literacy, encouraging students to critically evaluate the feedback they receive.

Additionally, Master of Teaching (Visual Arts and Design) students are using a combination of generative AI tools such as ChatGPT and DeepDream to learn about AI for creative and speculative writing exercises as well as developing designed objects and image-based works. Through this exploration, students are asked to reflect on authorship, consent, ethics, and creative practices.

ChatGPT has also been used by students in the Melbourne Law School to reflect on writing styles, reasoning skills and the use of case law.

Opportunities in research

Beyond the considerations already touched on above – enabling different ways of interacting with source materials, and using such tools to prompt novel ways of improving graduate researchers' reflective practices, generative AI presents a range of opportunities for research. For example:

- It could be used as a research tool if generative AI is integrated within library search tools.
- Generative AI could assist researchers when communicating to broader audiences, using non-academic language. For example, a researcher could use these tools to draft easily digestible newsletter text.
- Generative AI could help researchers undertake their research by developing code, generating faux test cases or interview questions, helping with the design of surveys, or testing methodologies. Already such use has been demonstrated in several disciplines.

Risk and challenges

While there are potential benefits, student use of generative AI raises new challenges to academic and research integrity and requires Universities to carefully calibrate their policy response. Such challenges are also posed by staff use of generative AI, requiring Universities to address what is acceptable practice for staff in using the technology for core activities, including setting assessments, marking, and providing feedback. Universities must be transparent in their policy decisions if they are to retain the trust of both the public and their student bodies.

Challenges to teaching, learning and assessment

The process of preparing material for assessment is an important part of students' education and university experience. It allows students to apply their learning and to develop analytical, communication, and presentation skills which are highly valued by employers. It also provides a means for both students and teachers to evaluate a student's progress, measure students' knowledge and skills, verify learning outcomes, and certify levels of achievement.

Generative AI tools have the potential to undermine the assessment process by quickly and cheaply producing novel outputs that are responses to a variety of prompts and are difficult to distinguish from student-produced work. While early tools had limitations in the datasets they used or the format of prompt they could respond to, these limitations are swiftly disappearing as the technology advances.

The University of Melbourne's *Student Academic Integrity Policy* states that all work submitted by an individual student must be their own.² As such, the use of generative AI tools to produce work that is subsequently submitted for assessment is considered academic misconduct, unless permission to use such tools has been given by academic staff and their use has been acknowledged by the student.³

As noted above, students are not banned from using these tools but must acknowledge material that is not theirs. A key difficulty in abiding by and enforcing this seemingly simple rule is the increasing ubiquity of generative AI tools within other standard software programs, and drawing a clear line between the levels of assistance that may be provided by, for example, an integrated thesaurus or grammar checker powered by generative AI.

Use of AI detection tools

In April 2023, Turnitin launched a tool that identifies material that may have been written by generative AI software. This tool is currently only available in staff view, a setting which cannot be changed by the University.⁴ While the validity and accuracy of the tool is still developing, the University is able to participate in testing and refining the tool through this early implementation period.

Staff have been advised that a high score within this tool is not to be taken as evidence of academic misconduct but may provide a flag for further investigation. Students have been advised that they may be asked to discuss or explain components of assessment tasks to assist with identifying potential breaches of academic integrity, and that this, in itself, does not constitute an accusation of academic misconduct. This is in line with existing practices around assessment integrity and within current policy.

In addition to working closely with Turnitin and the sector on understanding the tool and its accuracy, the University is running its own tests to understand the rate of false positives and false negatives and anticipates sharing this information with the community as it becomes available. The University's initial tests have suggested that these tools may be more likely to flag false positives where human authors use simple, predictable, or consistent word choices and sentence structures. The University has flagged these issues with Turnitin. Separately, the reliability of detection tools may vary as new, more sophisticated LLMs are developed. This highlights the importance of staff not relying solely on detection tools and the ongoing need to educate staff and students about the capabilities and limitations of generative AI.

It is worth noting that current evidence suggests that the use of generative AI for assessment and learning tasks is not yet widespread among students. A small University of Melbourne study from Semester 1 2023 found that almost half of all student respondents had not used generative AI. Additionally, fewer than one in ten students had used generative AI to create content submitted as all or part of an assessment, legitimately or otherwise. This is similar to Australian estimates of commercial contract cheating.⁵ However, user-friendly generative AI is still in its early days, and we can only expect usage to become more commonplace as time goes on and as the technology develops.

² The University of Melbourne (2022), *Student Academic Integrity Policy*, <https://policy.unimelb.edu.au/MPF1310/>

³ The University of Melbourne (2023), *Artificial intelligence tools and technologies*, <https://academicintegrity.unimelb.edu.au/plagiarism-and-collusion/artificial-intelligence-tools-and-technologies>

⁴ The University of Melbourne (2023), *Advice for students regarding Turnitin and AI writing detection*, <https://academicintegrity.unimelb.edu.au/plagiarism-and-collusion/artificial-intelligence-tools-and-technologies/advice-for-students-regarding-turnitin-and-ai-writing-detection>

⁵ Ziebell, N. & Skeat, J. (2023), *How is generative AI being used by university students and academics? Semester 1, 2023*, The University of Melbourne, p. 5

Challenges to research integrity

Research integrity refers to the issues under the *Australian Code for the Responsible Conduct of Research*. Academic publishers such as Nature and Elsevier have responded to the advent of generative AI tools with policy statements that restrict the use of AI-generated text in publications, prohibit the inclusion of AIs as authors, and set out guidelines for acknowledging the source.^{6 7}

The University of Melbourne strongly supports these statements and has released a similar statement.⁸ This clarifies that generative AI tools can only be used if material that is generated or substantially altered by these tools is acknowledged according to the University's policy and the *Australian Code for the Responsible Conduct of Research*.

The University's existing Authorship Policy requires those listed as authors to have made a significant intellectual or scholarly contribution to a research output and to willingly take responsibility for this contribution which excludes these tools from being named as authors.

Currently, the direct risks to publication integrity do not appear to be significant, but that may change. While current tools can be used to generate text that is intended for inclusion in a publication, such generated text is rarely at an appropriate academic level and is often wrong or absurd. Given that detection of submission of such material for publication would carry serious penalties, and that its use could be uncovered by a careful referee, this seems unlikely to be a common scenario, at least immediately. However, the generative tools are rapidly changing.

There is a much more substantial risk for Higher Degree by Research (HDR) candidates. For a novice writer of research material, the temptation to use generative tools is higher and the understanding of the consequences is lower. Development of communication skills is a key component of HDR study and mastery of fundamentals is essential before advanced capability can be achieved. Use of these tools may undermine acquisition of fundamentals. Writing is intimately linked with cognition and the ability to organise concepts into a coherent form. This includes not just the authoring of individual sentences and paragraphs but structuring of arguments, development of complex descriptions, and so on.

The process of grappling with how to precisely express concepts in written form is critical to development of them into research contributions. Likewise, assessment, review, and critique of writing is a key way in which supervisors mentor HDR candidates' development. This applies not just to theses but also to use of generative AI for material such as emails, literature reviews, proposals, and progress reports.

Finally, there are concerns around the disclosure of intellectual property (IP) or confidential information, as generative AI tools such as ChatGPT retain the prompts that are entered. The University has advised researchers that they should not share confidential information or information about an innovation in a generative AI prompt, as that may mean that the IP is no longer owned by the researcher or by the University.⁹ It is of note that the Australian Research Council recently released guidelines on the use of generative AI tools in ARC assessment processes.¹⁰ For example, the ARC has clarified that peer reviewers must not use generative AI as part of their assessment activities.

Privacy and bias

Generative AI raises risks of eroding information privacy and perpetrating bias.

As previously outlined, generative AI is trained on massive amounts of data scraped from the internet, such as books,

⁶ Nature (2023), *Artificial Intelligence (AI)*, <https://www.nature.com/nature-portfolio/editorial-policies/ai>

⁷ Elsevier (2023), *The use of AI and AI-assisted technologies in writing for Elsevier*, <https://www.elsevier.com/about/policies/publishing-ethics-books/the-use-of-ai-and-ai-assisted-technologies-in-writing-for-elsevier>

⁸ The University of Melbourne (2023), *Statement on research integrity and digital assistance tools*, <https://research.unimelb.edu.au/strengths/ethics/research-integrity-and-digital-assistance-tools>

⁹ The University of Melbourne (2023), *Graduate researchers and digital assistance tools*, <https://gradresearch.unimelb.edu.au/preparing-my-thesis/graduate-researchers-and-digital-assistance-tools>

¹⁰ Australian Research Council (2023), *Policy on Use of Generative Artificial Intelligence in the ARC's grant programs*, <https://www.arc.gov.au/sites/default/files/2023-07/Policy%20on%20Use%20of%20Generative%20Artificial%20Intelligence%20in%20the%20ARCs%20grants%20programs%202023.pdf>

articles, websites and posts. This data contains personal information that has been obtained without consent and without “contextual integrity”¹¹, raising serious privacy concerns. This data also includes text that is copyrighted or proprietary, creating issues around intellectual property when users re-appropriate this content in an educational or research context.

Other privacy and intellectual property risks emerge from the privacy policies of ChatGPT and other generative AI tools. ChatGPT saves users’ prompts, which may be reviewed by their AI trainers to improve their systems.¹² Noting this, OpenAI has advised its users not to share sensitive information in their conversations with ChatGPT. Universities will need to ensure staff and students are aware of these risks, to ensure that confidential information or information about innovations are not unwittingly shared, as this information may be provided by ChatGPT in responses to other people’s prompts. This risk may be reduced by new privacy settings, such as those added to a more recent version of ChatGPT, which preclude the use of data for training purposes.

Beyond privacy risks, Universities will need to educate staff and students about the biases present within generative AI models. Generative AI models reflect and perpetuate the biases of the data they are trained on. In the case of most generative AI models, this data is predominantly drawn from affluent Western societies and will therefore reflect the biases of people within those societies.

The representational and language-based biases this introduces are well documented. For example, Stable Diffusion, an AI-powered image generator, was found to largely produce images of white men when asked to create images of people in high-paying jobs. Conversely, it overwhelmingly produced images of women or people of colour when asked to create images of people in low-paying jobs. It found that it was worse than the real world, taking racial and gender biases to extremes.¹³ This could have significant consequences, perpetuating discrimination when generative AI is used in automated decision-making, including in policing and access to welfare. Indeed, these issues arise with any use of AI (generative or predictive) in automated decision-making. Such biases could also pose serious issues for the education sectors depending on their use. Concerns about bias in generative AI must become part of the education and training of students and staff generally, and in any direct use of generative AI.

Equity and access

Some students may be disadvantaged by increasing usage of generative AI. This includes students who lack reliable internet connections, access to digital technology, or even who are unable to afford the fees associated with subscriptions to state-of-the-art generative AI products. As noted by the Melbourne Social Equity Institute, approximately 2.5 million Australians had no internet connection in 2018. For many that did have access to the internet, this was limited by the quality or affordability of the service.¹⁴ As with existing technologies and technology-based tool provision, Universities should be active in ensuring equitable access of tools required for education and research.

There is a risk that the advent of generative AI will lead to greater stratification between those who are digitally literate and those who are not. This may be exacerbated by decisions by State Government’s to ban generative AI in public schools, for example, while some private schools begin to utilise generative AI in their education offerings.¹⁵ Education providers at all levels will need to carefully consider how these risks can be balanced.

¹¹ Contextual integrity refers to the principle that individuals’ information should not be revealed outside of the context in which it was originally produced. See further information:

<https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10/>

¹² OpenAI (2023), *What is ChatGPT?*, <https://help.openai.com/en/articles/6783457-what-is-chatgpt>

¹³ Nicoletti, L., Bass, D. (2023), *Human are biased. Generative AI is even worse*, Bloomberg <https://www.bloomberg.com/graphics/2023-generative-ai-bias>

¹⁴ Dulfer, N., Smith, C., van Holstein, E., Garner, A., Acosta Rueda, L., Rouse, L., Hamed, S., Cavanagh, K., and Ruppanner, L. (2022), *Understanding Digital Inequality: An analysis of unequal connectivity in Carlton Housing Estate, Melbourne, Victoria*, Australian Communications Consumer Action Network, p. 5

¹⁵ Carey, A. (2023), *Victoria banned ChatGPT in state schools. But Catholic schools took a surprising stance*, The Age, <https://www.theage.com.au/national/victoria/victoria-banned-chatgpt-in-state-schools-but-catholic-schools-embraced-a-surprising-stance-20230321-p5ctv3.html>

Future impacts

Generative AI tools are already having an impact on teaching and assessment practices, research, and operations in higher education, and this is expected to continue in the future. Generative AI will also influence the roles of educators, researchers, professional staff in educational institutions, who will need to refine their work practices as a result. Some of these likely impacts, particularly in the area of teaching, learning, and assessment, are discussed below.

Reconsidering assessment

TEQSA does not require Universities to use any particular forms of assessment. Instead, the *Higher Education Standards Framework (Threshold Standards) 2021* state that methods of assessment must be “consistent with the learning outcomes being assessed, are capable of confirming that all specified learning outcomes are achieved and that grades awarded reflect the level of student attainment”.¹⁶ The Threshold Standards also require that institutions take preventative action to mitigate foreseeable risks to academic integrity and provide students with guidance on what constitutes academic misconduct.¹⁷

However, the growing accessibility and capability of generative AI tools that can rapidly produce text, images, or other content in response to simple user prompts poses a new and significant threat to academic integrity. ChatGPT and other tools can easily produce seemingly impressive responses to prompts, making the user seem well-informed despite minimal effort and making it increasingly difficult to confirm that students have achieved learning outcomes. Universities and academics are currently considering how they can redesign assessments to reduce academic misconduct.

No assessment type can eliminate the risk of cheating. However, it is possible to design assessments that simultaneously facilitate the development and measurement of higher order complex learning outcomes and reduce vulnerabilities to misuse of AI. For example, students could be asked to complete small tasks in an iterative manner, showing evidence of the processes they have used to do so, and using platforms that record an audit trail of their progress. Students could also be asked to complete more hands-on, simulation-, object- or performance-based assessment tasks that allow them to demonstrate their learning and knowledge in ways other than writing. They may also be asked to complete more assessable tasks in-class (e.g., by incorporating assessment into teaching and learning activities). A large body of research can help guide the creation of assessments that students perceive as meaningful, authentic and relevant, and for which the motivation to cheat is reduced.

Ultimately, an effective response to the threat of generative AI will need to go beyond changes to individual assessment tasks and include consideration of assessment regimes at a whole-of-subject and program level. It will become increasingly difficult, to assure the integrity of every assessment task. Therefore, it may be necessary to shift the balance of assessment to include more ongoing, formative assessments with a smaller number of key, ‘gatekeeper’ summative assessments which are delivered in a format (e.g., individual presentation or viva) that provides high confidence of integrity. Many Universities were already shifting their approaches to assessment prior to the advent of ChatGPT to better support students’ learning. Concerns about generative AI may therefore simply accelerate shifts that were already underway.

As part of the University’s new *Advancing Students and Education Strategy*, we are committed to transforming our approach to assessment. This will include moving to more authentic, continuous forms of assessment that support student learning. As part of this work, we will build on the work of early adopters within our staff community to consider the ways in which generative AI can be embraced and integrated into our teaching. It is an excellent opportunity to think imaginatively about what we are really trying to assess and transform our assessment to focus on the learning process as much as the outcome of students teaching and learning experience.

Professional development of staff and education of students

Prior to using generative AI in teaching, learning and assessment, both students and academic staff need to develop AI literacies so that they can engage in safe and appropriate usage of these tools. Such understandings include the

¹⁶ *Higher Education Standards Framework (Threshold Standards) 2021* (Cth), s 1.4.3

¹⁷ *ibid.*, s 5.2

current strengths and limitations of the tools, the sources of the materials used to train the AI and any biases inherent within these, who has ownership of any data that are entered into the tool by the user, where those data are stored and how they will be used in the future. Both students and academic staff should also have clear understandings of how to use generative AI to produce work in ways that abides by research and academic integrity standards.

The Council of Europe distinguishes between two forms of AI literacy, which focus on the technological dimension (techniques and technologies of AI) and the human dimension (preparing for the impacts).¹⁸ They note that both need to be given equal billing and interwoven throughout, rather than seeing the human dimension as a “nice-to-have but inessential add-on.”¹⁹

It is also beneficial for academic staff to consider the appropriate use of generative AI tools in their teaching practice and research. For example, while using automated marking (a form of AI) to provide students with rapid feedback on a low stakes multiple choice quiz might be appropriate, efficient and useful, using generative AI or automation to provide feedback on *all* formative assessments would likely be detrimental for students. This is because it eliminates one of the few personalised and relational interactions that students have with their teachers as well as a valuable opportunity for teachers to monitor and support student learning, their learning development, and their improvements over time. It may also embed biases, particularly where it is used to provide feedback on more subjective forms of assessment. Further, automated marking technologies operate by recognising norms and may mark down work that is unusual, meaning that they can fail exceptional work.

Changing content material/education within programs to accommodate the use of AI in professions

Teaching, learning, and assessment activities at the University – as in the education sector worldwide – will need to adapt and respond to changes in technology and will need to do so in concert with industry to ensure our students are equipped to use generative AI tools as they move into their professions. The nature of work in many professions and industries will change because of new generative AI tools. Regardless of whether it is teaching, communication and marketing, software engineering, or legal practice, university educators will need to accommodate the role of generative AI in professional practice into their curriculum programs.

The disruption of professional practice in some industries has the potential to be significant; thus, educating for the future practice in these industries will need to evolve and accommodate this disruption. This will take time and effort. For instance, a review of intended learning outcomes will likely identify examples of outcomes that were historically relevant but are no longer fit-for-purpose in a world of generative AI. This may have the beneficial outcome of ‘raising the bar’ of expectation and achievement to include learning outcomes that are more focussed on uniquely human skills and attributes.

Impacts on the research workforce

An issue for research is that it takes place in a global rather than institutional environment. We cannot regulate what other researchers do and often have only limited mechanisms for calling them to account. The threats are at all stages in the research pipeline, including data and graph generation, writing of text, and writing of reviews on submitted papers. The risks of generative AI usage in research are unique in many ways to those faced in teaching, learning and assessment. Reliance on generative AI tools by researchers may lead to erroneous conclusions or unreliable research outcomes, for example. Universities will therefore need to ensure that education and training for staff and students on the use of generative AI includes tailored information and examples, noting the disparate impacts on different roles.

Recommendations to manage risks and seize opportunities

It is challenging to balance the potential of generative AI as an educational tool with the risks. It is worth noting that educators have been adapting to technological change for many years, so there are lessons to be learnt by reflecting

¹⁸ Holmes, W., et al. (2022), *Artificial Intelligence and Education: A critical view through the lens of human rights, democracy and the rule of law*, Council of Europe, p. 19

¹⁹ *ibid.*, p. 26

on past experiences. However, the pace of technological change and the scale and reach of impact across society does create other challenges from this new technology.

Governments may need to act in some form to help balance the need to protect users from misuse of generative AI with the need to encourage or facilitate innovation and allow the potential societal value of generative AI to flourish. There are existing approaches in consumer protection, privacy protection and other forms of regulation that may need to be updated, but this will be challenging with complex cross-jurisdictional issues and novel aspects of the emerging technologies.

The University acknowledges that the Department of Industry, Science and Resources is currently consulting on governance mechanisms for generative AI more broadly. However, it is worth considering specific actions that could be undertaken within the higher education sector.

Universities should continue to have autonomy over their institutions' policies and approaches to generative AI to ensure they are appropriate for their communities. In doing so, Universities and other education institutions will need to maintain transparency in their use of generative AI in core student- and staff-facing functions. However, the Government may wish to develop guidance for the sector on the use of generative AI, to aid internal considerations. For example, the Government could outline the various practical and ethical issues (including data privacy, biases, accountability, and reliability of generative AI) that institutions may wish to consider when developing and updating their internal policies.

Additionally, the Australian Government could work with the sector to establish communities of practice on generative AI in education between Universities, Government and industry partners. This could include conferences, workshops and forums where stakeholders exchange ideas, best practices, and lessons learned in generative AI in higher education. Participants could discuss and share alternative forms of assessment that reduce the risk of academic misconduct, creative ways to integrate generative AI into teaching and assessment, and emerging ethical issues.

The University would welcome TEQSA's continued guidance for the sector and sharing of useful resources. The current Threshold Standards are quite broad, requiring institutions to properly assess learning outcomes and mitigate risks to academic and research integrity. This means that updates to the Threshold Standards to reflect generative AI may not currently be necessary. However, this may change as the technology advances; TEQSA will need to monitor these risks and consider changes to the Standards in the future as required.

As highlighted in the Rapid Response Information Report, there is a risk to Australia associated with the concentration of generative AI resources within a small number of mainly US-based companies. Australia currently struggles to compete in the production or research of these technologies due to the significant computational and data requirements, as well as the need for highly skilled workers.²⁰ Overcoming these challenges will require significant investments from the Australian Government and industry in research and development of AI and related technologies. The Universities Accord panel may wish to consider this when proposing possible reforms to the Government.

Finally, there is a need for additional research into the role of generative AI specifically within education. Higher education in Australia is both foundational to the country's future and represents its fifth largest export industry. The enviable reputation of Australian higher education is founded on the quality and integrity of our educational approach. Generative AI puts this at risk. Given this, the Government might consider establishing a fund through the Australian Research Council or a Centre of Excellence to support leading-edge research in the development, use and impact of generative AI in education.

²⁰ Bell, G., Burgess, J., Thomas, J., and Sadiq, S. (2023), *Rapid Response Information Report: Generative AI - language models (LLMs) and multimodal foundation models (MFMs)*, Australian Council of Learned Academies, p.8

