



THE UNIVERSITY OF
SYDNEY

Professor Joanne Wright

Deputy Vice-Chancellor (Education)

Office of the Deputy Vice-Chancellor (Education)

14 July 2023

Ms Lisa Chester MP

Member for Bendigo

Chair

House of Representatives Standing Committee on Employment, Education and Training

Parliament House

Canberra ACT 2600

Submission lodged online through the Parliament's website

Dear Ms Chester,

Generative AI in education inquiry

The University of Sydney welcomes the Committee's inquiry into the use of generative artificial intelligence in the Australian education system and is pleased to provide the **attached** submission.

Our submission addresses the inquiry's six terms of reference and provides an overview of our approach to generative AI in our teaching, assessments, research and support activities.

This a fast-moving area of policy and practice, with similar discussions occurring globally – many of which we and our academic experts are connecting with through our educational, industry and research networks.

Next week, for example, we are hosting a symposium that will bring together learning and teaching leaders from more than 30 Australian and New Zealand higher education institutions to discuss best practices in curriculum development and assessment, in response to the rapid advancement of generative AI technologies.

We would be delighted to host members of the Committee for in-person or online meetings with the University's leadership and our experts working at the forefront of generative AI in our teaching, research and administration – if that would be of interest to the Committee.

Thank you once again for the opportunity to engage with the Committee regarding this important inquiry.

Yours sincerely,

Signature removed

Professor Joanne Wright

Deputy Vice Chancellor (Education)

Attachment The University of Sydney submission to the House of Representatives Standing Committee on Employment, Education and Training's inquiry into the use of generative artificial intelligence in the Australian education system, July 2023

University of Sydney submission to the House of Representatives Standing Committee on Employment, Education and Training's inquiry into the use of generative artificial intelligence in the Australian education system, July 2023

Contents

Executive summary and recommendations	1
Context	3
Responses to the inquiry's terms of reference	4
ToR 1 The strengths and benefits of generative AI tools in education	4
ToR 2 The future impact of generative AI tools on education.....	5
ToR 3 The risks and challenges for education and research posed by generative AI	6
ToR 4 Equity of access to AI tools for students	7
ToR 5 International and domestic responses to generative AI in education	8
ToR 6 Recommendations to manage the risks and seize the opportunities of generative AI ...	9
Endnotes	10
Appendices	11-26
A How Sydney academics are using generative AI this semester in class, March 2023	
B How AI can be used meaningfully by teachers and students in 2023, January 2023	
C ChatGPT is old news: How do we assess in the age of AI writing co-pilots? June 2023	
D Responding to generative AI for assessments in Semester 2, 2023	

Executive summary and recommendations

The University of Sydney welcomes the Committee's interest in the current and future implications of generative AI for Australia's education system, including the higher education sector. Generative AI is quickly becoming part of everyday life – changing the ways in which people work, interact with information and be creative. With ever-improving capabilities, generative AI will likely fundamentally change our students' lives and careers, as well as the knowledge, skills and strategies that our academics will need for success as educators and researchers.

Our vision for approaching generative AI is therefore based on an acceptance that this technology will become ubiquitous - augmenting and redefining our educational and research experiences, allowing our students and staff to focus on advancing human endeavours, and serving to enhance the quality and efficiency of our support operations. We seek to centre human agency in the application of AI by equipping our students and staff with the capabilities and dispositions necessary to thrive in an AI-augmented world, and by requiring the responsible, ethical and equitable use of generative AI across all aspects of our operations. Our vision includes nurturing an institution-wide culture of continuous learning, where every member of our community has equal access to the opportunity and resources to harness the potential of AI as they choose. We are also committed to raising levels of understanding about the opportunities, limitations and risks of generative AI in the wider public, through engagement with schools and communities, and by making our AI fluency and other relevant education and training programs as easily available as we possibly can.

In our response to **Term of Reference 1** and **ToR 2**, we provide a high-level overview of how we are encouraging and supporting our educators to embed generative AI safely and beneficially into their teaching, where they believe there is value in doing so. This is a fast-moving area, with significant experimentation and collaboration underway across the University, around Australia and globally.

To assist the Committee, our submission includes four appendices as examples of the type of guidance and support we are providing our educators. These include examples of beneficial use of the technology by our educators during Semester 1, 2023, as well as the guidance we have been providing around assessments.

Under **ToR 3**, we outline many of the key legal, ethical and other risks and challenges for teaching and research that we see generative AI posing for Australian academics and universities, and which we are focused on mitigating.

In response to **ToR 4**, we stress the critical importance of current and future generations of young Australians leaving secondary school with strong *foundational skills* (English language, literacy, numeracy, digital literacy and employability) if they are to participate effectively and meaningfully in society, post-school education and work. We also raise concerns about the potential for the *digital divide* in Australia to be widened on jurisdictional, socio-economic, cultural and geographic grounds. Achieving equality of access to generative AI for all Australian school students, regardless of backgrounds, is likely to require national policy leadership and potentially targeted funding or regulatory interventions.

Under **ToR 5**, we provide an overview of best practice developments regarding generative AI in higher education globally. In our assessment, the international perspective is overwhelmingly that the age of AI has arrived, reaching a level of sophistication and development that means it will significantly impact society at least in the longer-term, and so cannot be ignored.

Finally, under **ToR 6**, drawing on the content for our responses to the other terms of reference, we make the following five recommendations for the Committee to consider putting to the Australian Government:

Recommendation 1

Lead the establishment of a national cross-sector representative 'AI in Education' body and knowledge-sharing hub – to guide the safe, beneficial and equitable use of generative AI in Australian education, and to develop, test and showcase outstanding innovation in the productive and responsible use of the technology in the Australian school, vocational and higher education sectors.

Recommendation 2

Lead the development, in consultation with cross-sector representatives (see recommendation 1), of an initial national set of minimum standards to maximise consistency in the safe and equitable development and use of generative AI in the Australian school, vocational and higher education sectors.

Recommendation 3

Commit, for the next 3-5 years, to prioritising funding to support research and development work into the safe and effective use of generative AI to improve evidence-based teaching, learning and assessment in the Australian school, vocational and higher education sectors.

Recommendation 4

Identify policy and program opportunities to encourage and support sustainable collaborations between educators in schools serving disadvantaged communities and their colleagues in Australian universities with relevant expertise and interests, to help address the *digital divide*, specifically in relation to access to and use of generative AI tools.

Recommendation 5

Lead the development of a national set of principles and guidance materials targeted at research institutions to support and educate researchers (including research students) in the use of emerging technologies, including generative AI, in the research process. Development, for example, of a specific supporting guide for the [Australian Code for the Responsible Conduct of Research \(2018\)](#) addressing the benefits and risks of emerging technologies in research would enable all Australian organisations conducting research, to identify and consistently manage use of these technologies in their research and training programs.

Context

Defining 'generative AI'

For the purposes of this submission, we define *generative Artificial Intelligence* (Gen-AI hereafter) as a rapidly evolving class of computer algorithms able to create *digital content* – including text, images, video, music and computer code. They work by deriving patterns from large sets of training data that become encoded into predictive mathematical *models*, a process commonly referred to as 'learning'. Generative AI models do not keep a copy of the data they were trained on, but rather generate novel content entirely from the patterns they encode. People can then use *interfaces* like ChatGPT or MidJourney to input *prompts* – typically instructions in plain language – to make Gen-AI models produce new content.

The role of a university education in the age of generative AI

Globally, there is an enormous amount of discussion currently underway about the implications of Gen-AI for teaching and learning in our schools, universities and other tertiary education institutions.¹ This includes commentary, reflection and debate about the continuing relevance and value of longer-form qualifications offered by universities and of how best to prepare school leavers and university graduates for success in a world that is evolving rapidly due to advances in technology like Gen-AI and other forces including globalisation and climate change.² While written well before the public release of the ChatGPT interface and with future school education in mind, these words from our Vice-Chancellor, Professor Mark Scott AO, remain highly relevant to our current context:

While we cannot predict the future and the skills requirements of the employers of the future, we do know the types of learners that we want to develop... students who are critical and reflective, open to a lifetime of learning and relearning, who are comfortable with change and have empathy and a global outlook. This demands that we take a broader perspective about what we judge a good education to be...³

At the University of Sydney, we have put much research and thought into designing our curricula, learning outcomes and assessments to ensure that our students graduate with the following qualities: **deep disciplinary expertise; critical thinking and problem-solving; effective communication; information and digital literacy; inventiveness; cultural competence; interdisciplinary effectiveness (teamwork); and integrated professional, ethical and personal identities.**⁴

We agree with the Vice-Chancellor and others that we cannot be certain about the precise skills that people will need to operate effectively in their communities and societies in the future – to have successful and rewarding careers, and meaningful lives. Notwithstanding Gen-AI's remarkable and rapidly advancing capabilities, based on the currently available evidence and research, we have determined that we are right to continue focusing on developing the above qualities in our students, while also ensuring they are adept with the current Gen-AI tools and are well-placed to adapt and upskill as lifelong learners throughout their careers.⁵

The University of Sydney's approach to generative AI to date

Following OpenAI's public release of the ChatGPT interface in November 2022, our initial focus was on quickly determining and communicating our policies regarding the appropriate use of Gen-AI in our teaching, on ensuring the integrity of our academic assessments, and on supporting our staff and students to engage with the technology safely, responsibly and productively.

We have invested heavily in developing resources and training to assist our staff in both managing and making the most of Gen-AI in their teaching and assessments (unlike some other institutions who reacted with restrictions and prohibitions). We have consulted widely with staff, have included students in our AI in Education Working Group alongside educational and research experts, have mandatory academic integrity training in place for all students, and communicate with them regularly about our requirements, including the appropriate use of Gen-AI and other tools.⁶ We have also worked closely with students to develop resources that democratise information about Gen-AI so that all students can be upskilled.⁷

Alongside this education-focused activity, we have been working to establish an initial comprehensive institution-wide AI Strategy and governance framework that will set the direction for our journey with Gen-AI across our core educational and research activities, as well as the entirety of our supporting operations. Consistent with [UNESCO's guidance for education policy makers](#), the Australian Government's [Eight AI Ethics Principles](#) and the 'better community outcomes' principles underpinning the [NSW Government's AI Assurance Framework](#) our approach, once settled, will prioritise the following principles:

- human benefit, agency, expertise and accountability
- balancing the pursuit of Gen-AI's potential benefits for individuals, society and the environment with careful consideration of the risks
- the safety and wellbeing of individuals, communities and society
- equitable access
- ethical applications across all aspects of the University's work
- a supportive and collaborative approach to exploring AI use and governance
- normalising human-AI collaborative work
- privacy, information and data security, and transparency.

We are currently progressing the development of this strategy through consultations with our staff, students and external stakeholders, informed by a green paper. We are holding a national symposium in Sydney on 20 July 2023 to bring together learning and teaching leaders from over 30 Australian and New Zealand higher education institutions to discuss best practices in moving curriculum and professional development forwards in the face of Gen-AI. Finally, we are proactively considering the potential for Gen-AI applications to improve aspects of administrative support services. For example, we are in the early stages of running a proof-of-concept pilot to develop a plain English chatbot to help staff and students navigate the University's policy framework. We are using this pilot to inform the development of our overarching governance approach to Gen-AI.

Responses to the inquiry's terms of reference

ToR 1 The strengths and benefits of generative AI tools for children, students, educators and systems and the ways in which they can be used to improve education outcomes.

We have limited our response to this ToR to providing examples of how our educators and students have used ChatGPT and other AI tools beneficially in Semester 1, 2023. Case studies and other information elaborating on some of these use-cases are included in **Appendices A and B**.

Our students have told us clearly that we have a responsibility to teach them how to use Gen-AI tools properly. As one of our undergraduate students said in a consultation forum held about Gen-AI earlier this year:

"The traditional ways of learning... are changing. But this technology is now our present and the future, so we need our teachers to prepare us for it. A lot of us still do not fully understand this technology either and we need to be taught how to use it productively for ourselves, for tertiary education, and in greater society through our careers."

The current Gen-AI Large Language Models (LLMs) are known to produce flawed responses, including biases and inaccuracies, but they are rapidly improving. Meanwhile, the use of deepfake video and AI-generated virtual film locations is set to become more prevalent in the Australian and international film, TV and Subscription Video on Demand (SVOD) production industries. We need our students to be equipped with the knowledge to identify errors or weaknesses in the outputs produced by Gen-AI and to become masters in the effective use of these tools.

In the future, the ability of our students and graduates to judge the quality of their own work and to critically assess the outputs of their Gen-AI co-pilots and other models may become one of the most important qualities we need to develop in them. Future iterations of Gen-AI models will likely provide text, image and sound content that mimic human creative outputs to even higher levels. This will only increase our responsibility to adapt and evolve our approach to teaching, learning and student assessment in a Gen-AI world, so that our graduates continue to be sought after for their deep disciplinary knowledge, creativity, critical thinking and human-centric qualities.⁸

For educators, who are under ever-increasing workload pressures, Gen-AI has the potential to accelerate aspects of their roles. For example, Gen-AI can help educators propose new ideas for assessments, draft assessment rubrics, or generate practice questions for students. These activities are workload-intensive for an unassisted educator, but Gen-AI speeds up the process while still including the human in the loop. This frees up educators' time so that they can spend it on the aspects of their roles where the human value-add is more critical, such as student consultations, class time and personalised feedback.

Ten examples of beneficial use of generative AI by our educators

- Improving understanding of the technology, its limitations and appropriate use by students.
- Stimulating open class discussions using Gen-AI.
- Critiquing Gen-AI responses with students.
- Helping students research.
- Personalising teaching and assessment.
- Composing exemplars for critique with students.
- Generating discussion prompts.
- Drafting lesson plans.
- Drafting quiz questions.
- Designing draft marking rubrics.

Ten examples of beneficial use of generative AI by our students

- Conducting research using Gen-AI to gather examples and pointers for deeper investigation.
- Learning by teaching the Gen-AI about new concepts and receiving feedback from the tools.
- Asking the Gen-AI questions and critiquing its responses.
- Overcoming writer's block.
- Brainstorming ideas to receive feedback.
- Developing structure for essays and other forms of communication.
- Harnessing Gen-AI film and image production tools to extend creative ideas.
- Helping to create new code, images, music.
- Exploring diverse perspectives on a topic.
- Analysing texts with a research assistant.

See **Appendices A** and **B** for elaborations on some of these examples.

ToR 2 The future impact generative AI tools will have on teaching and assessment practices in all education sectors, the role of educators, and the education workforce generally.

Again, we limit our response to the university sector, its educators and our assessment practices in response to rapid advances in the capability and availability of Gen-AI tools. Regarding assessments, the verbatim use of Gen-AI tools to produce text or other outputs (images, video, code, designs) which a student then submits as their own, is not appropriate, and our updated [Academic Integrity Policy](#) reflects this. As indicated in our response to ToR 1, to ensure the relevance of our curriculum and the authenticity and integrity of our assessments, we have started to bring AI tools into the classroom and into our assessments, where our educational experts believe that doing so is appropriate and improves learning outcomes.

We have a responsibility to students, the community, and regulatory bodies⁹ to ensure that our students are meeting their learning outcomes. To combine the need to assure learning outcomes and embrace human-AI collaboration for assessment, our advice adopts a 'two lane' approach. **Lane 1** includes secure assessments predominantly in a live, supervised setting, designed to be as authentic as possible. **Lane 2** includes setting assignments where students are encouraged and taught how to collaborate with Gen-AI productively and responsibly, focusing on assessing the *process* of learning as well as the *product* of that learning.

See **Appendices C** and **D** for more examples of the types of guidance we are sharing with our educators to empower them in relation to their approach to assessment tasks following the advent of Gen-AI.

Some of our suggestions include:

- Adapting to Gen-AI, rather than trying to ban or outrun the technology.
- Rediscovering what it means to be human and assessing these skills and attributes.
- Refocusing on the desired student learning outcomes.
- Assessing the learning process as well as the product.

- Co-creating outputs with Gen-AI.
- Evaluating outputs co-created with Gen-AI.
- Asking students to describe and reflect on their use of Gen-AI and the lessons learnt.

ToR 3 The risks and challenges presented by generative AI tools, including in ensuring their safe and ethical use and in promoting ongoing academic and research integrity.

Some of the key risks and challenges of AI for academia we are focused on mitigating include:

In teaching

- Ensuring the safe use of Gen-AI by students and educators.
- Raising levels of awareness and understanding among teachers and students about the strengths, limitations and risks of Gen-AI.
- Overcoming perceptions that we can out-design or out-run Gen-AI by designing creative or authentic assessments that are 'AI-proof'.
- Ensuring staff and students understand what biases in Gen-AI training data can mean for information discovery, analysis and conclusion-drawing.
- Ensuring students do not develop an over-reliance on Gen-AI, nor a fear of the technology or reluctance to learn about it and engage with it.
- Preventing the inappropriate use of Gen-AI in assessments.
- Building students' ability to separate human from AI content.
- Reliance, at any level, on currently available AI detection software, which are inherently flawed.

In research

- Privacy breaches of study participants' personally identifiable information, through release of this information to Gen-AI tools, or the release of de-identified data in a way that allows for reidentification.
- Privacy breaches of researchers/collaborators' personally identifiable information, through submission of this confidential information to Gen-AI tools.
- Loss of control of Intellectual Property or Copyright, through the Gen-AI tools reproducing content without attribution.
- Breaches of contractual obligations, including commercial contracts, data sharing agreements and material transfer agreements.
- Non-compliance with the terms of funding agreements.
- Inappropriate management of First Nations data.
- Disclosure of sensitive environmental/ecological information.
- Breach of research codes, Human Research Ethics Committee (HREC) approvals (e.g. Australian Code for the Conduct of Responsible Research).
- Breach of research funding councils' specific policies regarding Gen-AI.

Other key risks for research integrity centre around lack of accountability and oversight when using Gen-AI to produce research outputs, specifically:

- Inability to attribute authorship where Gen-AI has contributed content from unknown sources.
- Inability to reproduce or replicate consistent results when using Gen-AI to provide content for research findings – results will vary at each attempt making verification of research findings impossible.
- Inability to establish and evaluate accuracy and validity of content generated by Gen-AI – probabilistic models that generate data may be biased, or produce 'hallucinations' of inaccurate content, which can be hard to identify.
- Inability to have contributory IP acknowledged by others if research outputs are incorporated into data that becomes available to the Gen-AI model, and is used elsewhere. No guarantee that licences on open data, code or images are applied (e.g. Creative Commons).
- Accountability and trust – personal accountability by individual researchers underpins all research. Use of Gen-AI can undermine this if its capabilities and the contribution that it makes, including risks, are not understood by researchers.

Additional potential risks and impacts for universities and their staff include:

- National security and foreign interference risks, arising from the public release through Gen-AI interfaces, of confidential data or research information, or from increasingly sophisticated cybersecurity attacks enabled by Gen-AI.
- Human rights and modern slavery law breaches, arising from use of certain Gen-AI tools developed without transparency and potentially through questionable labour practices.
- Researchers and research students who speak English as a second language may face increased pressure and risks, as they may be more likely to use LLMs that can improve written text, e.g. to help write a thesis or journal paper, in which case risk of IP loss is high.
- Some journal publishers have introduced rules restricting use of Gen-AI in submitted manuscripts. This varies between disciplines and publications, but may prevent some researchers from being able to publish in high impact journals.
- Detection rates of Gen-AI use are currently poor. This may have research integrity impacts in the future when new detection tools are applied to previously published research where Gen-AI may have been used to generate content. Detection tools may give erroneous results, and valid research outputs may be questioned in error.
- Reviewers of grants and awards using Gen-AI inappropriately in the process of evaluating proposals to circumvent human thought and insight.¹⁰

ToR 4 How cohorts of children, students and families experiencing disadvantage can access the benefits of AI.

We will leave it for experts in early childhood and school education to advise the Committee on evidence-based steps that could be taken to bridge the ‘digital divide’ that may prevent many Australian children and their families from accessing the benefits of Gen-AI.¹¹

Ensuring equity of access to generative AI systems

We do need to stress, however, the increasing importance of young people, regardless of their socio-economic backgrounds, leaving school with strong foundational skills (English language, literacy, numeracy, digital literacy and employability skills). Technologies like Gen-AI are likely to continue the transformation of work in the Australian and global economy that has been underway since the industrial revolution, further reducing the share of jobs that require routine skills, while potentially increasing demand for people with higher-order cognitive skills.¹²

In the near future, all Australian students will need to leave school with a deep understanding, not only of how to use Gen-AI, but of how it works, as part of their foundational skills needed if they are to be critical, creative users of AI, and to participate fully in an AI-augmented society and economy. Strong foundational skills are critical to an individual’s ability to engage successfully in tertiary-level studies and life-long learning, to avoid obsolescence and find new occupations, if necessary, that require retraining or upskilling.¹³ Foundational skills also underpin an individual’s capacity for critical analysis and to discern truthful and accurate information from the intentionally fake or otherwise erroneous information that increasingly sophisticated Gen-AI models will be able to produce.

However, according to the Grattan Institute, data from the OECD show that the performance of Australian school students in reading and maths is going backwards, whether compared to past performance or international benchmarks. Worryingly, the learning gap between students from advantaged and disadvantaged backgrounds more than doubles between Year 3 and Year 9.¹⁴ Meanwhile, according to the Productivity Commission, one in five Australian adults have poor foundational skills, leaving them unable to perform basic tasks, more likely to be unemployed, disengaged from education or training and on lower incomes.¹⁵ Australia can improve how students and families experiencing disadvantage access the benefits of Gen-AI by focusing policy and funding to maximise the prospects of all Australians completing secondary school with strong foundational skills and AI-literacy.

There is a significant risk that a new *digital divide* could widen quickly in terms of Australian school students’ access to and familiarisation with Gen-AI, because of resourcing disparities between schools and inconsistent policies pursued by different states and school systems regarding the technology.¹⁶ Access to increasingly powerful and specialised Gen-AI tools is unlikely to be free in the future. The danger is real that students from more affluent and educated families will have access to much more powerful (fee-for-service) applications.

For students enrolling with the University of Sydney, we recognise that we have an obligation to ensure that access to any Gen-AI technology accepted as appropriate in a field of study is accessible to all students in a cohort on an equitable basis. Our capacity to deliver on this commitment will, of course, depend on the commercial arrangements that the companies behind different Gen-AI models will apply for public not-for-profit Australian universities. This is one area where the government may be able to assist, whether through helping to reduce the subscription costs for providers, or with funding to help provide targeted Gen-AI access and learning support to students from underrepresented groups.

Building levels of generative AI fluency across the Australian community

Australia's universities, including the University of Sydney, are home to significant research and teaching expertise about all aspects of Gen-AI. Meanwhile, levels of understanding about Gen-AI and trust in the technology across the wider community are currently reported to be relatively low.¹⁷ The University of Sydney Business School offers a highly-regarded [AI Fluency](#) microcredential, targeted at executives in business, government and other organisations. The University of Sydney and many other Australian universities offer a wide range of educational offerings about different aspects of Gen-AI, including a [resource we have co-designed with students](#) to democratise access to information about Gen-AI for students. We see significant potential for Australian universities to work with governments, school systems, the teaching profession, communities and industry to rapidly expand levels of AI fluency, technical capability and teaching skills. Upskilling pre- and serving teachers about Gen-AI should be prioritised as they will shape students' understanding of the technology and attitudes to it through approaches to teaching and assessment. This is an area where recommendations from the Committee about the importance of such partnership approaches would be timely and helpful.

ToR 5 International and domestic practices and policies in response to the increased use of generative AI tools in education, including examples of best practice implementation, independent evaluation of outcomes, and lessons applicable to the Australian context.

The international policy and principles landscape around Gen-AI and education involves a balance between productive engagement and careful acknowledgement of the technology's limitations and ethical implications. In May 2023, the US Department of Education's Office of Educational Technology published a new policy report¹⁸ emphasising that Gen-AI enables new approaches to teaching, assessment, and feedback, while being cognisant of the risks especially those around data privacy and security. In the UK, the Russell Group of research-intensive universities recently released five guiding principles¹⁹, stating that all of its member universities will support the development of AI literacies, adapt teaching and assessment, work collaboratively, upskill students, and ensure integrity is upheld. We expect the Group of Eight universities in Australia to release a similar set of principles soon. UNESCO has also published a guide for educators about ChatGPT and similar Gen-AI interfaces, citing many ways that Gen-AI can improve teaching and learning, whilst highlighting its challenges such as bias, academic integrity, privacy, and commercialisation.²⁰

Overwhelmingly, the international perspective is that the age of Gen-AI is upon us, and it has reached a level of sophistication and development that it can make a real impact on society.²¹ Although initial reactions tended to be negative and centred around worries for academic integrity, attitudes to Gen-AI in education are shifting rapidly, indicating an accelerating desire for adopting Gen-AI with more people seeing that the benefits outweigh the drawbacks.²² However, institutional policies tend to be lagging²³, and the perspective of individual educators range from fanatical to adversarial.

With an international policy landscape that is cautiously optimistic about Gen-AI in education, the focus has been on: raising educators' awareness of Gen-AI; demystifying the technology; rapidly upskilling educators in the use of Gen-AI; redesigning educational practices (especially around assessment); and working with students as key partners in developing responses.

Students are critical stakeholders and in March 2023, a survey by Tyton Partners showed that a larger proportion of students than educators used Gen-AI tools, and that students are likely to use it even if it is prohibited.²⁴ The Jisc UK National Centre for AI has engaged well with students through focus groups and forums, yielding insights about how students see Gen-AI in their education and future careers.²⁵ At the University of Sydney, we have also elevated the student voice around Gen-AI in education, highlighting many positive ways that students use Gen-AI to learn, how institutions have a responsibility to help them develop Gen-AI literacies, and acknowledging its limitations.²⁶

Educator professional learning is also critical. Institutions, faculties, and other educator communities run countless workshops, seminars, and talks on Gen-AI in education, and there are key resources available

to support educators. Monash University's TeachHQ resource on Gen-AI and teaching and learning is widely applauded.²⁷ Our own work at the University of Sydney has also been used nationally and internationally, such as our 'Prompt engineering for educators'²⁸ and 'Prompt engineering for students'²⁹ series, our resources on assessment redesign³⁰, and our workshops on using Gen-AI productively and responsibly. As we discover more about this new general-purpose technology³¹, the key practice we must all adopt as educational institutions and individuals, is one of curiosity and openness to collaboration with one another.

ToR 6 Recommendations to manage the risks, seize the opportunities, and guide the potential development of generative AI tools including in the area of standards.

Drawing on the information and ideas we have raised in our responses to the other terms of reference, we recommend that the Committee consider recommending to the Australian Government that it:

Recommendation 1

Lead the establishment of a national cross-sector representative 'AI in Education' body and knowledge-sharing hub – to guide the safe, beneficial and equitable use of generative AI in Australian education, and to develop, test and showcase outstanding innovation in the productive and responsible use of the technology in the Australian school, vocational and higher education sectors.

Recommendation 2

Lead the development, in consultation with cross-sector representatives (see recommendation 1), of an initial national set of minimum standards to maximise consistency in the safe and equitable development and use of generative AI in the Australian school, vocational and higher education sectors.

Recommendation 3

Commit, for the next 3-5 years, to prioritising funding to support research and development work into the safe and effective use of generative AI to improve evidence-based teaching, learning, and assessment in the Australian school, vocational and higher education sectors.

Recommendation 4

Identify policy and program opportunities to encourage and support sustainable collaborations between educators in schools serving disadvantaged communities and their colleagues in Australian universities with relevant expertise and interests, to help address the *digital divide*, specifically in relation to access to and use of generative AI tools.

Recommendation 5

Lead the development of a national set of principles and guidance materials targeted at research institutions to support and educate researchers (including research students) in the use of emerging technologies, including generative AI, in the research process. Development, for example, of a specific supporting guide for the [Australian Code for the Responsible Conduct of Research \(2018\)](#) addressing the benefits and risks of emerging technologies in research would enable all Australian organisations conducting research, to identify and consistently manage use of these technologies in their research and training programs.

Ends/

Endnotes

Appendices

- Appendix A** How Sydney academics are using generative AI this semester in class, March 2023
- Appendix B** How AI can be used meaningfully by teachers and students in 2023, January 2023
- Appendix C** ChatGPT is old news: How do we assess in the age of AI writing co-pilots? June 2023
- Appendix D** Responding to generative AI for assessments in Semester 2, 2023

Endnotes

- ¹ <https://scholar.google.com.au>
- ² Loble, L., & Hawcroft, A. (2022) [Shaping AI and Edtech to Tackle Australia's Learning Divide](#), University of Technology Sydney, [How Important Is a College Degree Compared to Experience? \(hbr.org\)](#)
- ³ Scott M. Preparing Today's Students for Tomorrow's World, in Loble L. et. al. Future Frontiers: Education for an AI World, NSW Department of Education (2018) p.98
- ⁴ [Graduate qualities - The University of Sydney](#)
- ⁵ See for example: [The Skills Imperative 2035: An analysis of the demand for skills in the labour market in 2035 - NFER](#), Productivity Commission (2023) [5-year Productivity Inquiry: From Learning to Growth, Inquiry Report – volume 8](#), University of Sydney (2023) [Reimagining Australian tertiary education: 12 nation-building ideas, Submission to the Australian Universities Accord](#), p.13
- ⁶ <https://www.sydney.edu.au/students/academic-integrity.html>
- ⁷ Our 'AI in Education' resource for students is openly accessible: <https://bit.ly/students-ai>
- ⁸ <https://www.sydney.edu.au/study/why-choose-sydney/employability-and-careers.html>
- ⁹ For example, the Higher Education Standards Framework (Threshold Standards) 2021 around confirming that specified learning outcomes are achieved;
https://www.legislation.gov.au/Details/F2022C00105#_Toc67664702
- ¹⁰ [ARC Policy on the Use of Generative Artificial Intelligence in the ARC Grant Programs \(7 July 2023\)](#)
- ¹¹ Here we note, for example, the valuable contribution to this aspect of the discussion made by Leslie Loble AM in, [Shaping AI and edtech to tackle Australia's learning divide](#), published in December 2022
- ¹² See Productivity Commission (2023) [5-year Productivity Inquiry: From Learning to Growth, Inquiry Report – volume 8](#), p.iv
- ¹³ <https://www.dewr.gov.au/foundation-skills>
- ¹⁴ Wood, D., (2021) [Thinking Big: Maintaining full employment and boosting productivity](#), presentation to the Australian Government's Jobs and Skills Summit, 1. Sept. Slide 19
- ¹⁵ Productivity Commission (2023) [5-year Productivity Inquiry: From Learning to Growth, Inquiry Report – volume 8](#), p.9
- ¹⁶ [NSW Department of Education Policy on AI in NSW public schools](#), [South Australia Department of Education Guidance on AI in public schools](#), [South Australia AI pilot](#)
- ¹⁷ KPMG and the Australian Information Industry Association (2023), [Navigating AI, Analysis and guidance on the use and adoption of AI](#)
- ¹⁸ <https://tech.ed.gov/ai/>
- ¹⁹ <https://russellgroup.ac.uk/news/new-principles-on-use-of-ai-in-education/>
- ²⁰ <https://unesdoc.unesco.org/ark:/48223/pf0000385146>
- ²¹ <https://www.gatesnotes.com/The-Age-of-AI-Has-Begun>
- ²² <https://er.educause.edu/articles/2023/4/educause-quickpoll-results-adopting-and-adapting-to-generative-ai-in-higher-ed-tech>
- ²³ <https://www.insidehighered.com/news/2023/03/22/gpt-4-here-most-faculty-lack-ai-policies>
- ²⁴ <https://tytonpartners.com/generative-ai-in-higher-education-from-fear-to-experimentation-embracing-ais-potential/>
- ²⁵ <https://nationalcentreforai.jiscinvolve.org/wp/2023/07/03/exploring-the-significance-of-generative-ai-for-students/>
- ²⁶ <https://educational-innovation.sydney.edu.au/teaching@sydney/student-staff-forums-on-generative-ai-at-sydney/>
- ²⁷ <https://www.monash.edu/learning-teaching/teachhq/Teaching-practices/artificial-intelligence>
- ²⁸ <https://educational-innovation.sydney.edu.au/teaching@sydney/prompt-engineering-for-educators-making-generative-ai-work-for-you/>
- ²⁹ <https://educational-innovation.sydney.edu.au/teaching@sydney/prompt-engineering-for-students-making-generative-ai-work-for-you/>
- ³⁰ <https://educational-innovation.sydney.edu.au/teaching@sydney/chatgpt-is-old-news-how-do-we-assess-in-the-age-of-ai-writing-co-pilots/>
- ³¹ <https://arxiv.org/abs/2303.10130>

Appendix A

Educational Excellence, Teaching tips, 7 March 2023

How Sydney academics are using generative AI this semester in class

by Danny Liu

Our approach to ChatGPT and other generative artificial intelligence (AI) at the University of Sydney is one of [productive and ethical engagement](#). Whilst recognising the need to safeguard academic integrity, we want to help [prepare students to be ethical leaders](#) in a future where AI provides new and exciting possibilities.

In the first part of this series, we spoke with academics who are experimenting with ways to help students better understand generative AI and use it to power-up their research and discussions. These are all low-workload but high-impact practices that you could try in your own classes. Helping students explore and understand ChatGPT

Holding open class discussions

[Kim Weatherall](#), a professor in the Sydney Law School, coordinates senior elective units on Law and Technology. “The focus of the units are on the relationship between law and technology: how we regulate new and emerging technologies; how law can or should respond when societies, economies, human capabilities and behaviour are changed by developments in new technologies,” she says. ChatGPT is the perfect contemporary example for her students to wrestle with. Kim is allowing her students to use it during the process of putting together a research essay, but they aren’t allowed to submit text written by these tools.

To gear up her students to think critically about ChatGPT, she ran a class discussion where students were active and engaged. “I demonstrated it live in class: we asked it to explain some law & tech issues; then write a joke about AI, then write a haiku about AI, then write an essay outline for us. It was an eye opener for some of the students I think. I then put a series of questions to small group discussion before we talked as a whole class,” she reflects. To encourage open discussion, **she shared a collaborative document with her class and invited students to note their comments into the document**. After demonstrating ChatGPT, she posed questions around benefits, risks, and ethics – and she has [shared her discussion prompts for others to use](#).

Kim is engaging with these AIs in class because it is part of her students’ future:

By allowing, and encouraging them to use the tool and explore its capabilities, students will develop a deeper understanding of it, which will serve them both in the course and in future careers if they pursue practice in this area.

Critiquing ChatGPT’s responses together

[Charlie Warren](#), an associate professor in the Faculty of Science, wanted to use ChatGPT as a “trendy hook” to highlight that it was still important to learn to think critically. In his third-year undergraduate unit on Applied Plant Function, the class activity he ran “served as a nice real-world demonstration of ChatGPT’s limitations and, by inference, the very real need for human cognition.” He started with an overview of how scientists would usually research a topic – evaluating and then synthesising multiple reputable sources.

He then demonstrated ChatGPT responding to a trivial short answer exam question, about which Charlie quips, “I knew ChatGPT would get wrong because the majority of responses on the internet are misconceptions”. ChatGPT was then asked to summarise and write to research questions about trees and carbon dioxide sequestration. Students then worked in groups to evaluate and improve ChatGPT’s responses, which Charlie had pre-generated before class so students didn’t need to create their own accounts. **Groups were asked to mark up Word documents and comment on logical problems, misinformation, shallow analyses, and other issues**. Finally, students were asked to reflect on their experience with ChatGPT and its utility as a research tool. Charlie has [made available his class activity](#) for others to adapt.

Like Kim, Charlie's goal was to encourage his students to approach ChatGPT realistically:

It's an emerging technology that students are using. We cannot pretend it doesn't exist or attempt a ban. It seems reasonable that we help students explore its limitations.

Encouragingly, the initial exposure helped students to become more comfortable with experimentation. Charlie reflects that his class "seemed to enjoy exploring the limitations of ChatGPT as a research tool, while also recognising how it could be useful (e.g. for initial research, setting out a framework, and/or overcoming writer's block)," and that some even went further to try and engineer better prompts.

Take a minute to ask ChatGPT something tricky/tough

- Choose a topic that you have a sizeable amount of knowledge in (anything)
- Assess how well ChatGPT has given the answer
- Could you prompt it again to give you a better answer?
- Could you have changed the first prompt to give you that better answer to begin with? What could you have done that would be different?
- Pair up with a partner and discuss what you found, before we have a class discussion

The University of Sydney

Page 107

Hamish's class activity

In a similar way [Hamish Fernando](#), a lecturer from the Faculty of Engineering, asked his students to examine ChatGPT's responses in his second-year units on AI, Data, and Society in Health. He encourages students to use generative AI for creative assignment and writing ideas, and as a friendly programming assistant. To help his students build critical familiarity with ChatGPT, he ran **a class activity where students generated a response from the AI on a topic in which they had substantial knowledge, and worked together to critique its quality** and attempt to improve the prompt to generate better responses.

Helping students research

[Carolyn Stott](#), a senior lecturer in the Faculty of Arts and Social Sciences, started engaging with ChatGPT herself because she was fearful. "I was initially scared of the ramifications for assessment!" she reflects. Now, she uses ChatGPT in her senior undergraduate French unit to assist students' research processes and bolster their understanding of language. "In my French units I have openly discussed ChatGPT as one piece of software – among many others – that might contribute to their language learning, if used appropriately. I have encouraged students to trial it in class (in French) as part of formative learning activities e.g. research a definition of a particular literary genre and share it with the class or research the significance of a particular date or group of actors in the Algerian War," she says.

While these activities using ChatGPT have produced good responses, she also includes activities that ChatGPT struggles with. "I have also shown them what happens when ChatGPT is asked to come up with literary analysis of a number of texts", she says, noting that ChatGPT does not do this well because some of the texts she uses are not well known. She plans to use it to teach grammar, "by encouraging students to ask it to provide simple grammatical explanations of topics they don't understand."

Carolyn wants her students to have a realistic approach to ChatGPT. "Discussing openly the pros and cons of this tool is important," she notes. "Students can see how it can help them in formative tasks and perhaps at the beginning of a summative task – like they might use Wikipedia, for example. They can also understand the ethical issues around its use and hopefully make the appropriate decisions."

Spurring class discussion

[JJ de Vries Robbé](#) teaches in International Financial Transactions and International Sustainable Finance, postgraduate intensive units in the Sydney Law School. To help students ease into the subjects, he includes short pre-class exercises, some with ChatGPT. He also uses ChatGPT as a means to kickstart discussions in class. **“My recommendation to students is to use ChatGPT for what it is: as a brainstorming buddy.** At the same time I caution on believing the prompts – do not simply accept, but challenge what’s been put before you. This applies to ChatGPT, but of course to many things in life” he says.

Personalising teaching

[Tatjana Seizova-Cajic](#), a senior lecturer from the Faculty of Medicine and Health, uses ChatGPT to teach quantitative study design to undergraduate, Honours, and postgraduate students in different faculties. With diverse cohorts, it’s hard to find examples of interest to all, let alone discuss topics of special interest. But with ChatGPT, she now has “a teaching assistant who knows something about each field”.

“After I introduce a concept or study design, [students request an example in their field of interest from ChatGPT](#)” she says. “We discuss and evaluate ChatGPT output. It’s important to have the time to do it.” **ChatGPT allows Tatjana’s students to quickly contextualise a new concept or methodology within their preferred topic area**, making it come alive and helping to motivate their engagement with new concepts.

ChatGPT can also help to break awkward silences, by giving students a starting point for discussion and critique. “Once I described a simple study in the lecture, and asked if a causal conclusion was justified. Silence. I encouraged them to ask ChatGPT by typing in the scenario and question. The answer was clear and several students volunteered their version,” she reflects. This also provides an opportunity to discuss ChatGPT’s limitations and errors. “We have had fun in class with some references it confabulated.”

Tatjana finds that her students are better engaged through using ChatGPT, and using it actually encourages them to learn more:

It develops critical thinking because students soon realise that confidently presented text could be wrong. It does not replace thorough sourcing of information or the need to develop a body of knowledge.

Her students seem to agree – early feedback suggests that they approach ChatGPT as a generalised information source and see its limitations (such as factuality), whilst recognising that it can be more useful for activities like summarisation and translation.

As we lead our students through this new frontier of teaching and learning, our role as teachers is to model transparent, ethical, and productive engagement with generative AI. Tatjana has some wise words for educators everywhere: “Sure, some forms of assessments have to change, but our main role is to teach and I don’t want to put assessment before teaching. Students know that ChatGPT cannot replace their own understanding of material. **We cannot and should not try to hold them back from using AI out of concern that their professional growth will suffer. I think it’s quite the opposite: it will suffer if we don’t help them use AI.**”

Appendix B

Assessment, educational integrity, inclusive teaching, teaching tips, 20 January 2023

How AI can be used meaningfully by teachers and students in 2023

by [Danny Liu](#), [Erica Ho](#), [Ruth Weeks](#) and [Adam Bridgeman](#)

There has been [a lot of hype](#) over the last few weeks about students using ChatGPT, a conversational artificial intelligence (AI) that can write human-sounding responses to prompts. Some of the main concerns are that students might use ChatGPT to answer assignment or exam questions, without developing their own writing or critical thinking skills, or disciplinary knowledge. [In a previous article](#), we covered how GPT works, and some important concepts that help us appreciate what this AI does and doesn't do. Importantly, we emphasised that ChatGPT is not a knowledge model but a *language* model, albeit a powerful one. Here, we share some ideas about how we might productively engage with ChatGPT, a form of 'generative' AI, to make higher education learning, teaching, and assessment better in many ways – for both teachers and students. [Check out the other resources in our AI and Education at Sydney collection.](#)

For teachers

We're all thinking the same thing – not *another* disruption that we need to deal with. First emergency remote teaching because of COVID, then myriad systems changes, and now artificial intelligence? But let us share a few ways that AI can actually make your teaching easier and improve students' learning.

Writing quiz questions and feedback

ChatGPT prompt and response for writing MCQs.

A perennial challenge is providing students with meaningful formative questions on which they can practise, and with accompanying feedback that helps them learn from mistakes. We've [promoted Peerwise before](#) for its ability to engage students in deeper understanding by generating their own questions and feedback, whilst creating a question bank for practice. But if there's little space in your curriculum for that, consider using ChatGPT to generate practice quiz questions with feedback, and load these into a Canvas quiz. Given the inaccuracies and tendency for verbosity [we highlighted in the earlier article](#), **ensure that you check the accuracy and difficulty of the questions and feedback that ChatGPT generates**, and the alignment to your learning outcomes. Your students will thank you for the extra practice.

Example prompt for ChatGPT: Write three college-level multiple choice questions that target key student misconceptions around how natural selection works. Include feedback for students.

Designing a draft marking rubric

ChatGPT prompt and response around generating a marking rubric.

A marking rubric is a powerful table that helps you, your markers, and your students better understand what is expected against marking criteria at different standards of performance. A good rubric will reduce the workload associated with marking and engage your students with the feedback they receive. Each marking criterion is a row in the table, and standards (high distinction to fail) are presented as columns, with a descriptor for each criterion/standard. Rubrics, however, can be difficult to write from scratch. If you provide ChatGPT with a short description of your assessment, and what you'd like students to be assessed on, **it can generate a draft rubric that you can then finesse**.

Example prompt for ChatGPT: Design a marking rubric for a postgraduate assessment that asks students to apply their knowledge of the global financial crisis to a more contemporary economic challenge. The rubric needs to assess students on their use of literature, their analysis of the underlying causes of the GFC, and apply it creatively to a contemporary challenge. Please provide standards for each criterion from high distinction, distinction, credit, pass, and fail.

Composing exemplars for critique

ChatGPT prompt and response to generate a short assessment exemplar.

In addition to rubrics, another powerful tool for helping students better understand expected standards is to provide exemplars. Exemplars are samples of work, often presented with commentary at different levels of performance, to demonstrate and explain the expectations. Exemplars can be hard to come by, especially for a new assessment. Instead of composing an exemplar from scratch, consider asking ChatGPT to generate a few for you, and then amend these drafts based on your

expectations and standards. **You might want to be up front that the exemplars were produced by ChatGPT** to highlight its current linguistic and conceptual limitations, and to model the transparency we expect of students when they use it.

Example prompt for ChatGPT: Create a sample student paper that analyses the underlying causes of the GFC. This paper should achieve a credit standard, demonstrating only a basic understanding and identifying only one or two key causes.

(It's worth noting that when ChatGPT responded to the above text, it also commented that "It's worth to notice [sic] that this sample student paper demonstrates a basic understanding of the causes of GFC, it identifies only one or two key causes and lacks a critical evaluation of them, it also doesn't connect the causes to the consequences of the crisis.").

Generating discussion prompts for use in class

ChatGPT prompt and response to generate some discussion ideas.

Because ChatGPT is trained on a huge diet of written content available on the internet, it has linguistic coverage of a vast array of topic areas, opinions, and other text. Based on the examples it 'learned' during training, it can give statistically-likely responses to any textual prompt. Although this can sometimes mean that the AI generates inaccurate or biased content, it can provide text that represents human-like opinions and perspectives fairly well. You can leverage this to spark discussion in class, without spending too much time racking your brain for these. Depending on your discipline, you might ask ChatGPT to generate some discussion prompts, edit them based on your expertise and experience, and use them in class. Again, you should consider being up front about the source of the prompts, to **illustrate that ChatGPT reflects the biases present in the human-written sources upon which it was trained**.

Example prompt for ChatGPT: Bill is a 21 year old football player with an active social life and busy study calendar. He lives on campus in student accommodation. What are five reasons that Bill may give to their physiotherapist for not following a treatment plan for a tendon injury?

Writing lesson plans

ChatGPT prompt and response for composing a lesson plan.

If you've inherited a new class and haven't been given a plan for the lessons, or perhaps you'd like some ideas on how you might include more active learning in your classes, consider using ChatGPT to draft a lesson plan for you. Because ChatGPT was trained on dialogue so that **it can answer follow-up questions in a conversational way, once it generates an initial lesson plan, you could ask it to refine it**, such as by explaining how a particular activity may work, or suggesting how you might make an activity more motivating for students. Importantly, remember to apply your own pedagogical and content knowledge to evaluate and finesse the AI-generated suggestions.

Example prompt for ChatGPT: Write a lesson plan for a 50 minute college tutorial on the implications of artificial intelligence for future work. Include activities where students work in groups to critique different perspectives. Student groups should also investigate different industries that might be impacted by AI.

For students

Most of the [hysteria](#) around ChatGPT has been focused around its potential for misuse by students to complete [written assignments](#) and [exams](#). This has led many academics and universities to [refocus on good assessment design](#) and consider how to productively engage their students with AI to help them prepare for [the inevitable future where AI augments human intelligence](#). There may also be important roles for ChatGPT to [support students living with communication disabilities](#). Here are a few ideas about how you might use AI in your units to help students learn.

Learning by teaching

After ChatGPT was told that its solution is incorrect, it invited the user to teach it how to solve the equation correctly.

One of the strengths of language models is their ability to respond to prompts given by users in natural language and generate human-like text in response. As both the prompt (input from users) and response (output from AI) are in natural language, the user interface is friendly and models trained on dialogue such as ChatGPT can be potentially used as conversational agents that enable students to [act as a tutor and learn by teaching](#). When learning a new concept, students can explain the concept to ChatGPT like explaining to another student. And ChatGPT is likely to be a friendly peer

because it is fine-tuned to behave in a kind and gentle way. Although sometimes it may still provide less-helpful outputs, the AI has been improving and adapting its behaviour through learning from human feedback.

Students can also ask ChatGPT questions and the potentially inaccurate answers generated by the AI may offer opportunities for students' learning. When interacting with ChatGPT, students should be reminded to **always critically review its answers and explain why they think the answers provided by the AI are correct, inaccurate, or incorrect**. Critical thinking is an ever-important skill in today's knowledge society and in the face of misinformation and fake news on the Internet. Through reviewing the AI's answers, challenging its answers via follow-up prompts, judging the correctness by looking for supporting evidence or explanations, and experiencing the incorrect or inaccurate answers presented in a confident and persuasive way (like the simple mathematical problem demonstrated in the earlier article), students may gain a deeper understanding of the importance of critical thinking and be more sensitive when viewing plausible information. In addition, since ChatGPT is able to admit its mistakes, students can inform and teach the AI after noticing its incorrect or inaccurate responses (as demonstrated in the screenshot to the right). Using ChatGPT could be seen as an interactive way of learning, and the processes of identifying and explaining the errors (similar to [studying erroneous examples](#)) and providing correct answers with explanations could facilitate learning.

Overcoming writer's block

Like us, students might struggle to start a written submission because they are not used to a particular style of writing, or are uncertain about expectations, or are anxious about their knowledge and abilities. You might **consider explicitly encouraging your students to use ChatGPT to draft some introductory lines, topic sentences, or other parts** of the written work to get them started whilst explicitly highlighting its limitations. (Because, let's face it, your students will probably be doing this anyway). In this way, and where appropriate, the AI may also be used to address language inequality and help those who are equally creative but with less-developed writing or language skills (e.g. non-native speakers) to express their ideas. It's important to set boundaries and expectations, and perhaps spend some class time demonstrating how they can engage with AI to help catalyse writing – more on this a bit later.

Discussing with a co-programmer

ChatGPT is also able to take in and generate computer code. Professional programmers [have suggested that the AI can be a great companion](#) and explains snippets of code, suggests improvements to code, and helps to explore alternative approaches. In and out of the classroom, you might consider suggesting to students to use ChatGPT to suggest explanations for unknown code snippets or generate a boilerplate function for performing a given task. **Again, it's important to remind students that GPT is a large language model, not a knowledge model**, so the explanations and code that it generates [may contain inaccuracies or other issues](#). AI won't replace programmers or programming classes (yet), so it's important that students understand that they still need to develop their own coding knowledge and skills in order to check ChatGPT's output including its biases.

Exploring diverse perspectives

In a similar approach to the suggestion to teachers above about generating discussion prompts for class, you could consider asking students to use ChatGPT to generate a set of different opinions or perspectives on a topic. Because ChatGPT has been trained on so much human text, it is likely able to generate text that closely mimics real human perspectives on a range of topics. Students **could then use this as stimulus material for discussion, critique, and analysis** – and every student could generate and use their own unique set of perspectives, which has benefits for individualising assessments.

For assessments

A discussion of ChatGPT in higher education wouldn't be complete without directly considering its impacts on assessment. We have covered some uses above where AI could support assessments, such as drafting rubrics, creating exemplars, and individualising assessments. It's clear that some students could, and indeed [probably already are](#), use it to complete assignments and exams. There are already examples of where this has been detected by markers mostly due to the poor quality of the answers (we'd [advise against AI-based detection of AI writing](#)). As teachers, our role is to

consider how assessment processes may motivate students to learn, instead of focusing on blocking and subversion and detection of AI. As the inimitable Maha Bali expressed on Twitter: Sector leaders from Deakin University, which has an active research group on higher education assessment and learning, [have asked](#) whether “we [can] really afford to expend precious assessment space on narrow tasks” in a race back to the ostensible safety of invigilated exams. Instead, what are authentic and contemporary approaches to assessment in a context where AI will only become more intelligent and prevalent?

Get creative and incorporate AI into assessment

Ryan Watkins from George Washington University has [proposed a compelling range of suggestions for assessments](#) that **not only mitigate the risk of AI but also encourage students to develop their AI literacy**. These include having students generate ChatGPT responses and critiquing, evaluating, and visibly revising them. For example, AI could generate a series of short paragraphs in response to an assessment prompt, and students need to present the generated text and evaluate the ‘facts’ and perspectives contained therein, providing an accurate commentary supported with key references.

Set boundaries and expectations

The Sentient Syllabus Project has published an [open-access set of ‘syllabus resources’ that teachers can adapt for their own classes](#), which aim to emphasise to students three principles: (i) AI alone cannot pass a unit; (ii) AI-generated material needs to be attributed; and (iii) AI use should be transparent. Their document contains a wide selection of well-considered text that you can adapt and use, that will **help to set boundaries and expectations for your students around how they can and should productively engage with AI**.

Consider a range of assessment types that are chunked

A common response to guarding against ChatGPT is to make students refer to in-class materials or discussions, or contemporary sources. **While well-intended, students could still generate large swathes of the assessment using ChatGPT and just amend it** to include these references. ChatGPT can even generate [convincing personal reflections](#), so even reflective writing is not a guarantee of protection. Instead, consider [breaking up larger assessments into chunks](#) that need students to visibly iterate on their work through outlines and drafts before a final submission, and include multimodal (e.g. audio, video, pictorial, or other multimedia) outputs. (That said, upcoming AI technology such as [VALL-E](#) can generate terrifyingly realistic spoken audio just using a three second sample of someone’s voice). [Oral assessments](#) such as vivas are also a possibility, and are considered [one of the best approaches](#) to ensure assessment security.

Be open with students, and each other

As we start semester 1, ChatGPT will likely be one of the main elephants in the (class)room. Anna Mills from the College of Marin has curated a [set of questions that teachers should be asking about generative AI](#). In fact, these are great questions to pose to students as well. The questions encourage us to consider how these AIs work and their benefits and limitations, what we value in learning, the implications for equity and inclusion and [bias](#), and more. **Having these open conversations with students and each other will encourage us all to be transparent (and honestly, it’s OK to be a little terrified) and work out how to appropriately engage with AI together**. We are educating students to be leaders in the world and so we need to acknowledge openly that AI is part of the world now and will live and evolve with us. Be open with students that we want them to use but also lead/disrespect the technology. It doesn’t short circuit their learning. It won’t replace them or us as long as we value critical thinking and education as something beyond a transaction.

As teachers and a university that are committed to developing students’ deep disciplinary expertise as well as key skills such as critical thinking, communication, and leadership, we have a responsibility to develop their (and our) AI literacy, to productively engage and work alongside this powerful technology that has the potential to make our work, study, and leisure more effective, efficient, and inclusive.

Appendix C

ChatGPT is old news: How do we assess in the age of AI writing co-pilots? 8 June 2023

by [Danny Liu](#) and [Adam Bridgeman](#)

In the seven months since ChatGPT was released, the world of generative AI has progressed at an almost impossible pace. Trying to keep up with weekly updates on its impact on education has been like drinking from a firehose. The next big thing on the generative AI front that we need to pay urgent attention to are **AI writing co-pilots that will be directly embedded into productivity suites like Microsoft Office**. These “co-pilots” are AI-powered assistants designed to assist you to generate content.

The companies behind the productivity tools that most of us and our students use, Microsoft and Google, are getting their versions of AI co-pilots ready for imminent release. [Microsoft 365 Copilot](#) and [Google Duet AI](#) will embed generative AI directly into software like Microsoft Word and Google Docs. **This means that we will all have access to text generation AI right from within the spaces where we write.** Notion, a popular app used by many students, already has an [AI-powered writing assistant](#).

Ethan Mollick, a pragmatic and thoughtful educator with AI from the Wharton School in the US, describes this as a “[crisis of meaning](#)”:

I don't see a way that we avoid dealing with this coming storm. So many important processes assume that the amount and quality of our written output is a useful measure of thoughtfulness, effort, and time. If we don't start thinking about this shift in meaning now... we are likely to face a series of crises of meaning, as centuries-old approaches to filtering and signalling lose all of their value.

So, what happens to assessments now?

[Jason Lodge](#) from the University of Queensland and colleagues, and [Michael Webb](#) from Jisc UK's National Centre for AI, have written about the main options regarding assessments. Webb writes that we can either avoid it, try and outrun it, or adapt to AI. All of these responses are legitimate, and all have their short- and long-term benefits and drawbacks. Avoiding it includes reverting to oral or otherwise invigilated exams, which have stronger assessment security but higher costs. And, as assessment experts Liz Johnson, Helen Partridge & Phillip Dawson from Deakin University say, exams can have questionable authenticity. That said, there is certainly a place for secured assessments in order to assure learning outcomes are met – more on this towards the end of this piece.

Outrunning it involves trying to design assessments that AI has more difficulty completing – but **the risks are that our redesigns will only be temporarily effective as the pace of AI development accelerates and will make the assessment more inequitable for many of our students**. Modifying stimulus (e.g. using images in questions) has been suggested, but [GPT-4's](#) ability to parse images will be released to the public. Modifying the content of assessment is also a popular suggestion, such as connecting with personal events, writing reflections, or linking to class material – but recent work has shown that [GPT writes higher-quality reflections than humans](#), and a larger context window allows students to send it as much relevant class material as needed. Modifying outputs that students produce is also a common response, such as changing assessments to in-class presentations or multimodal outputs – but rapid advances in [voice cloning AIs](#) and [video](#) and [image generation](#) AIs make outrunning AI almost futile. Besides, students could just (and have, to excellent effect) use AI to generate an in-class presentation. This approach also has the potential to exacerbate achievement gaps for those students skilled in the use of AI and those who are not.

Adapting to it means we need to rethink how we assess – as Lodge writes, this is a more effective, longer-term solution, but also much harder. The imminent arrival of AI writing co-pilots makes this even more important – **students will have AI asking to be invited into their writing process right from the start, through an innocuous prompt to us humans like “Help me write” (Google) or “Describe what you'd like to write” (Microsoft)**. With this capability in place within the mainstream productivity tools also makes it inescapable – it will become specious and even meaningless for us to “ban” their use.

“Assess the process”, “use authentic assessment” – but how exactly?

Changing assessments to assess the process and not the product has been a growing retort over the last 6 months in response to generative AI. Lodge puts this elegantly in [a recent post](#):

*While generative AI can increasingly reproduce or even surpass human performance in the production of certain artefacts, it cannot replicate the human learning journey, with all its accompanying challenges, discoveries, and moments of insight. It can simulate this journey but not replicate it. **The ability to trace this journey, through the assessment of learning processes, ensures the ongoing relevance and integrity of assessment in a way that a focus on outputs cannot.*** (emphasis added)

Lodge mentions “reflection activities, e-portfolios, nested tasks, peer collaborations, and other approaches” as effective process-based assessments, but acknowledges that these “often do not scale easily or leave too much room for threats to academic integrity”. This is especially the case in the coming era of AI writing co-pilots.

And even before generative AI burst on the scene, there have been calls to recalibrate higher education towards more authentic assessment. Through a systematic review, Verónica Villarroel and colleagues determined there were three main dimensions to authentic assessment:

1. Realism (having a realistic task and realistic context)
2. Cognitive challenge
3. Evaluative judgement (developing students’ ability to judge quality)

Over the years, we’ve collected a few [examples of authentic assessment on Teaching@Sydney](#), and the University of Queensland has a [handy searchable assessment database with examples of authentic assessments](#).

But again, in a time when generative AI is so in-your-face and part of the way we [will] work, what is ‘authentic’, and what is ‘process’ – and how do we actually assess? Are our assessment approaches built to be *effective* or to be *efficient*?

Rediscovering what it means to be human (and assessing this)

Going back to first principles, the Higher Education Standards Framework 2021 legislation reminds us that assessments need to assure that learning outcomes have been met, and that the way that we assess needs to be consistent with these learning outcomes. Specifically, [Part A section 1.4 clauses 3 and 4](#) say:


3. Methods of assessment are 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.
4. On completion of a course of study, students have demonstrated the learning outcomes specified for the course of study, whether assessed at unit level, course level, or in combination.

There have been calls for a revisit to learning outcomes in the age of AI – the Australian Academic Integrity Network has [posted on the TEQSA website](#) that “unit and course learning outcomes, assessment tasks and marking criteria may require review to incorporate the ethical use of generative AI”. **Given the close and dynamic interplay between learning outcomes and assessments, the increasing presence of AI in common productivity tools suggests we may need to rethink learning outcomes as well.**

What do we want our students to know and be able to do when they leave our units and our courses? At Sydney, we want our graduates to have the [skills and knowledge to adapt and thrive in a changing world, improving the wellbeing of our communities and society](#). But what does this actually mean in practice, in the context of an AI-infused world? Cecilia Chan and colleagues from the University of Hong Kong have [asked a similar question, from the perspective of educators](#). They suggest that there are key human aspects that AI can never replace, such as cultural sensitivity, resilience, relationships, curiosity, critical thinking, teamwork, innovation, ethics, civic engagement, leadership, etc. If these sound familiar, it’s because they’re also embedded in [Sydney’s graduate qualities](#) for students.

So perhaps a contemporary set of learning outcomes (and assessments) needs to address these human elements. The ability to be able to judge the quality of their own work and that of their co-pilot may be the most important quality we need to develop. Sure, generative AI can create text and other outputs that mimic these human qualities, but that enhances our prerogative to adapt assessments appropriately to an AI world.

Designing assessments in an era where generative AI is inescapable



Generative AI and assessments (click to enlarge)

Consider this diagram, which we recently presented in the closing keynote of a UK conference on 'Education fit for the future' and had useful conversations on via Twitter afterwards. Much of the early dialogue around generative AI has been around its (mis)use in the bottom right quadrant – high AI contribution and low human contribution. **As generative AI becomes more inescapable, we need to consider where along the top half of the diagram each of our assessments sit** – and they will necessarily sit at different places, depending on the year level, learning outcomes, and other factors.

Leon Furze helpfully characterises this as an 'AI assessment scale', where educators need to consider whether and how AI helps students in assessments. We've adapted his scale in the diagram below, to emphasise that AI can also help educators in the design and delivery of assessments.



Where AI might be useful in assessment (click to expand)

What are some approaches to assessments that might work in the age of AI writing co-pilots, considering all we've discussed about process, authenticity, and learning outcomes? How might students work productively and responsibly in the top-right quadrant?

Co-creating an output with AI

There are many ways that industry and community groups are already and could potentially leverage generative AI, including content generation, customer insights, software development, media generation, document summarisation, knowledge organisation, and more. What are the steps that we and our students can take through assessments in this context? Here is a strawman for consideration:

1. In your field or discipline, **consider the kinds of authentic tasks and authentic contexts that your students may eventually find themselves meeting**. This may be producing an annual report for stakeholders, a design proposal for a firm, a policy briefing for government, marketing copy for a client, a scientific report for a journal, a pitch deck for investors, an environmental impact analysis for a local council, a proposal for a museum exhibition, a clinical management protocol for a hospital, a citizen science project outline for the community, etc. At this stage, you don't need to necessarily think about AI. (AI could help you, as an educator, dream up some of these ideas – see how).
2. Ensure the assessment is aligned to your learning outcomes. If it doesn't align, you may need to either refine your learning outcomes, or try a different assessment.
3. Draft a marking rubric for the assessment, which is aligned to the learning outcomes. **As part of this rubric, you may want to consider incorporating some of the very human qualities** encapsulated in the Graduate Qualities as rubric criteria (rows). For example, you

may want students to develop and demonstrate information and digital literacy skills, as well as inventiveness. As part of the rubric, consider adding criteria that will allow you to evaluate students' approach to co-writing – see the next section.

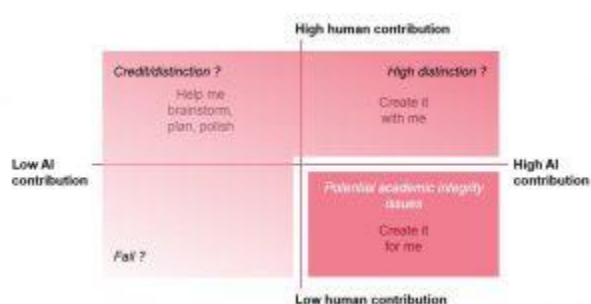
4. In the process of starting the assessment, **have students design a good prompt for generative AI, and document this process.** There are many considerations when designing a good prompt, and the quality of the initial prompt (and follow-up prompts) will determine the quality of the AI output. To design a good prompt that leads to a meaningful output, students will need to clearly specify the task, provide relevant context and examples, and even split complex tasks into subtasks or steps. **You can already see that a meaningful engagement with generative AI relies heavily on students having developed some disciplinary expertise.** It's important for students to document this as it demonstrates their learning process, critical thinking, and the cognitive challenges they overcome.
5. Once an acceptable draft has been generated by the AI co-pilot, have students critique and enhance this, and document what they do, how they do it, and why. The aim here is for students to improve the AI's outputs by analysing its accuracy, judging its quality, integrating relevant scholarly sources (the AI is unlikely to yield any reliable sources) – all elements of developing evaluative judgement. Again, **it's important that students document what they do, how, and why, as it serves as a tool for you to assess the process.**
6. If there is space in your curriculum, implement a draft submission at this stage. Students can get feedback and critique from peers, who have also gone through a similar process of using an AI co-pilot.
7. Since we are interested in assessing the process and not (as much) the outcome, at this stage have students submit the amended draft, their documentation (initial prompt and reasoning, and their analysis and improvement steps), their reflections on peer feedback, and their proposed changes going forwards. **Assess these as the primary (i.e. most heavily weighted) component of the overall assessment grade – for ideas on how, see the next section.**
8. Have students improve the output again, documenting the process as before.
9. Students submit the final output. This, and their documentation, is assessed, but has a lower weighting because we are interested in assessing the process, not the outcome.

Evaluating co-creation with AI

An integral part of the suggested model above is to be able to evaluate the process of making something *with* AI. Some of the following ideas may be useful criteria to include in the marking rubric, alongside other criteria. Ensure that the criteria that you include in your rubric align with your learning outcomes and the Graduate Qualities.

- **AI prompt design that demonstrates disciplinary expertise:** how thoughtfully the student has designed the prompt(s) for AI and considered the complexity and clarity of prompts. High distinction could include a well-structured and deep understanding of disciplinary concepts that are demonstrated in effective prompt design.
- **Critical evaluation of AI suggestions:** how effectively the student evaluates and utilises AI suggestions, as in whether they simply adopt AI-generated content or make conscious choices about what to include. High distinction could include critically making nuanced and evidence-based decisions about what to accept, modify, or reject.
- **Revision process:** how the student has revised AI suggestions and demonstrated their critical thinking skills and disciplinary expertise. High distinction could include an insightful and critical reflection on where AI generated content needed improvement, and why. It could also include a demonstrably significant improvement in the quality of the work.
- **Information and digital literacy:** how the student has evaluated AI-generated content through relevant scholarly sources to enhance the rigour and reliability of the output. High distinction could include the integration of high-quality sources that are appropriately critiqued.
- **Documentation and reflection on the co-creation process:** how the student has recorded appropriate decisions and interactions with the AI co-pilot, and analysed the strengths, weaknesses, and future improvements to these interactions. High distinction could include a clear and ethical articulation of decisions and their reasoning, deep insight into the role of AI in the co-creation process, and suggestions for future practice.
- **Ethical considerations:** students' awareness of the reliability, biases, and other limitations of AI generated content. High distinction could indicate a strong understanding of these issues, with suggestions for mitigating potential problems.

These are by no means an exhaustive or prescriptive list. You will need to determine the most effective way to assess and evaluate these skills for your own teaching context.



Could the human-AI collaboration elevate student performance? (click to enlarge)

In contexts where collaboration with AI is acceptable, authentic, and productive, it may be conceivable that a student who creates something *with* AI may learn more, better hone their critical thinking and information and digital literacy skills, and produce a better artefact. To very loosely paraphrase [Ethan Mollick](#), **students may be hurting themselves by not using AI, if AI-enabled writing is of a higher quality**. Certainly, we need to be fiercely conscious of many issues of collaborating with AI, not the least because of [anchoring bias](#) where the first (in this context, AI-generated) piece of information clouds our judgement.

Assuring learning outcomes

Webb's strategies of either [avoiding](#), [trying to outrun](#), or [adapting to AI](#) are not necessarily stark choices and our final position will no doubt be a mixture of all three.

We want to know our students (and our future bridge builders and dentists) know what they are doing: our purpose and social licence is based on our ability to assure that our graduates have met their learning outcomes. This will no doubt involve some secure summative tasks in degrees and programs which exclude technology to ensure program level outcomes are met. Given the inevitable play off between authenticity, security and cost, we need to consider the responsibilities of each unit versus those of the program. With generative AI baked into our productivity tools, this may involve appropriate methods of assurance in these exceptional 'high stakes' assessments. As Cath Ellis from UNSW [suggests](#), **perhaps we need to stop agonising about over-securing every single assessment**. Assessments where students collaborate with AI may help them learn critical skills and develop deeper disciplinary expertise, whilst assessments that are highly secured may help us (and them) additionally assure learning outcomes are met.

If we are honest, our present [assessment regime](#) reflects our need to deliver *efficiency*, reflecting the workload implications of large enrolments and tight deadlines. Assessment for and of learning is not the only consideration that a unit coordinator tasked with returning results to a deadline and within a limited budget needs to consider. Whilst generative AI may offer some efficiencies in [designing and delivering assessments](#), it also requires us to consider where and what to assess.

As the costs, security, and reliability of 'traditional' assessments soar, is this a time to think about assessing *where* learning happens – such as in the laboratory, in the design studio, and in the tutorial? **At the moment, assessment weightings and hence student behaviours favour products (notably assignments and exams) whereas formative learning, development, and evaluative judgement occurs in the tasks in our active learning classes.**

Adapting to AI will take time and will probably require a mixture of student skill building, staff development, and a paradigm shift in our understanding of what assessment is for. For students, this will involve foundational digital literacy in using AI (such as generic prompt writing) but this will eventually need to be a discipline-specific skill and part of their knowledge creation methodologies. For staff, we need to rapidly and widely build awareness and support the innovations of early adopters.

The ultimate aim of a Sydney education will perhaps not be changed by the generative AI revolution. Indeed, it should become more real and relevant. Our assessments should require, encourage, rank, and reward the ability to use evaluative judgement on the quality, reliability and relevance of the resources our students use and the outputs they produce.

What now, and what next?

- Consider the implications of AI co-pilots in your disciplinary context. Consider what impacts this will have on assessments, particularly written assessments.
- Given the HESF requirements, ensure that select assessments are appropriately secured so that your unit and program can demonstrate student achievement of learning outcomes. For other assessments, consider the above guidance on adapting to and embedding AI.
- Find out through industry and community groups and other stakeholders how AI is being used (or not) in your field, and where it is effective (or not). Use this to inform your decisions about teaching and assessment.
- A working group at the University is drafting a green paper towards a potential University-wide AI strategy. This will seek to help guide the University along a path of ethical, equitable, responsible, and productive use of AI in all facets of the University's work. Wider consultations on this will be opened soon.

Appendix D

Responding to generative AI for assessments in Semester 2, 2023

by [Danny Liu](#), Associate Professor at University of Sydney, July 4, 2023

[Generative AI](#) is improving at a rapid pace and is increasingly integrated into software that we use daily. Together with folk across the University of Sydney, we put together a quick guide to help coordinators ensure that assessment changes have longevity, even as AI progresses. The guide also emphasises the opportunities for us to reflect on the human side of teaching and learning, help students become better prepared for AI-augmented careers, and encourage the ethical, accountable, and transparent use of AI.

Here is a summary of the contents of the guide; you are welcome to [download the full guide](#). Note that the guide has been produced for use within the University of Sydney, but it has broader applicability across different institutions.

Diagnosing and addressing impact

Discover the capabilities. Run your assessment prompt(s) through generative AI, preferably using the more powerful models available (such as [GPT-4](#) for text generation, Adobe Firefly for image generation, etc). Ensure you play around with the prompts to properly explore AI's potential capabilities (see Appendix 1 in the [full guide](#)).

Evaluate student motivations. Students are more likely to undertake assessment [with integrity if they feel](#):

- Autonomy: having real choice about topic and mode, and seeing how the assessment *meaningfully* connects with their life and career.
- Competence: being supported to build confidence and skills gradually.
- Relatedness: feeling connected to teachers and peers and that they matter.

Balance assurance of learning and use of AI. All units need to have assessments that assure student attainment of learning outcomes – this is likely best performed in live, supervised settings. It's also critical to ensure we help students use AI productively and responsibly, which can be done by redesigning other assessments to address appropriate learning outcomes. The table on the next page provides guidance for balancing these.

Reduce the perceived workload and pressure. Assessments with clear instructions and criteria, have meaningful and appropriate challenge, that provide sufficient time for completion, and which help students develop confidence in their abilities (e.g. through structured drafts and feedback) will lead to more positive academic integrity outcomes.

Decide and communicate. It's important to differentiate AI use for *learning*, and AI use for *assessment*. Use this guide to determine what level of AI usage is appropriate in context, and clearly communicate this to students – including situations where AI use is not appropriate. Examples of appropriate use and wording about how to acknowledge AI use are provided in the [full guide](#). The [full guide](#) has specific advice for different assessment types including take-home assignments, small continuous assessments, groupwork assessments, and exams and tests.

Redesigning assessments

Some level of assessment redesign is required across almost every unit to both manage the risk of generative AI *and* provide students with opportunities to engage with it productively and responsibly. In a world where AI is inescapable, assessments should both assure learning in secure settings, *and* adapt to the reality of AI in other settings, as appropriate to each discipline.

The two-lane approach below emphasises balance between assurance, and human-AI collaboration. **The reality in any one unit will likely be a situation where some assessments lie in lane 1 in order to assure attainment of all learning outcomes, but most other assessments lie in lane 2.** Fundamentally, we want to develop students who are well-rounded and can contribute and lead effectively in authentic, contemporary environments (which will include AI), and also be assured of

their learning. Therefore in this context, **it is important to privilege lane 2 assessments with a higher weighting than lane 1 assessments.**

Lane 1 – Examples of assured ‘assessment of learning’

- In-class contemporaneous assessment e.g. skills-based assessments run during tutorials or workshops
- Viva voces or other interactive oral assessments
- Live simulation-based assessments
- Supervised on-campus exams and tests, used sparingly, designed to be authentic, and for assuring program rather than unit-level outcomes
-

Lane 2 – Examples of human-AI collaboration in ‘assessment as learning’

- Students use AI to suggest ideas, summarise resources, and generate outlines/structures for assessments. They provide the AI completions as an appendix to their submission.
- Students use AI-generated responses as part of their research and discovery process. They critically analyse the AI response against their other research. The AI completion and critique provided as part of the submission. Appendix 3 in the [full guide](#) provides suggestions on how to assess this.
- Students initiate the process of writing and use AI to help them iterate ideas, expression, opinions, analysis, etc. They document the process and reasoning behind their human-AI collaboration. The documented process demonstrates how the collaborative writing process has helped students think, find their voice, and learn. The documented process is graded and more heavily weighted than the artefact. Appendix 3 in the [full guide](#) provides suggestions on how to assess this.
- Students design prompts to have AI draft an authentic artefact (e.g. policy briefing, draft advice, pitch deck, marketing copy, impact analysis, etc) and improve upon it. They document the process and reasoning: initial prompt, improvements, sources, critiques. The documented process demonstrates learning, is graded, and is more heavily weighted than the artefact. More information. Appendix 3 in the [full guide](#) provides suggestions on how to assess this.

We do not foresee a viable middle ground between the two lanes. It needs to be assumed that any assessment outside lane 1 (i.e. that is un-secured) may (and likely will) involve the use of AI.

An example of assessments across both lanes

In this example, students need to apply marketing strategy concepts in real-world scenarios, demonstrate their communication skills, and evaluate the effectiveness of different marketing strategies.

The lane 2 assessment might involve students collaborating with AI such as Bing Chat (which is internet-connected) to perform market research and competitor analysis, and other AI such as Adobe Firefly for the visual elements of campaign design. Students document their interactions with the AI tools, including the AI's initial market research and analysis and their critique and fact-checking processes to evaluate the AI's outputs. Students also critique whether AI provided novel insights and whether it missed critical factors. This is then presented live in class. The grading of the assessment is more heavily weighted on the documented process of critical co-creation (see Appendix 2 and Appendix 3 of the [full guide](#)).

The corresponding lane 1 assessment might involve a live Q&A after the presentation, where students need to defend their research and analysis through targeted questions. This can be made to simulate real-world business meetings, and helps to assure that students have met their learning outcomes of applying marketing strategy concepts and evaluating effectiveness of marketing strategies. Another lane 1 assessment might involve giving students an unseen case study of a company that has recently launched a new product; in a live, supervised setting, they need to evaluate the effectiveness of the marketing strategy and propose areas of improvement.

Using AI as part of assessment

The [full guide](#) provides examples of uses of AI in assessment, as well as wording that could be presented to students to explain this. It also provides examples of how students might acknowledge the use of AI.

Some examples include:

- Generating ideas for assessment: You may use AI tools such as <ChatGPT, Bing Chat, and other generative AI> to <brainstorm ideas and approaches> for completing your assignment.
- Creating media for assessment: You may use AI tools such as <DALL-E, MidJourney, Stable Diffusion, Adobe Firefly, and other image generative AI> to generate <images> that you use as part of your submission.
- Providing feedback on work: You may use AI tools such as <ChatGPT, Bing Chat, and other text-to-text generative AI> to seek feedback on your written work.
- Searching literature: You may use AI tools such as <elicit.org, perplexity.ai, and researchrabb.it.ai> to find and summarise research articles. You then need to incorporate the scholarship yourself into your submission.

Detecting students' AI use

Software that purports to detect the use of AI-generated text is prone to false positives and false negatives. Research has suggested these tools are [not reliable](#) and may be [biased against](#) non-native English writers. In addition, ChatGPT does not 'know' whether it generated a piece of text – even if it may produce a convincing response when asked.

Our advice is that you should not submit students' work to AI detection software yourself – this is potentially a breach of student privacy and intellectual property.

Testing your own assignment with generative AI

It's important to test your assignment(s) against generative AI to gauge what kinds of outputs could be produced. When prompting generative AI, it's important to remember that better prompts will yield better results – don't stop at the first prompt/response and dismiss AI if its initial response is not impressive.

The [full guide](#) has a step-by-step guide on how to properly try out your assignment prompts on AI.

Suggestions for rubric criteria that target higher order thinking skills

A marking rubric is a tool that allows teachers, markers and students to form a shared understanding of the specific criteria and standards used to make academic judgements. A rubric directs students' work by providing descriptions of the standards at different levels of achievement. Appendix 2 of the [full guide](#) has some sample wording that you should adapt for your needs.

Note that the rubric suggestions attempt to privilege the more human elements of writing and composition as part of assessment design and grading. It is becoming trivially easy for AI to replicate these elements if prompted the right way. Therefore, changing your rubric should **not** be the only change you make to assessment in response to AI.

Suggestions for rubric criteria that privilege the process of human-AI collaboration in assessment rather than the product

Appendix 3 in the [full guide](#) provides rubric examples that can be adapted for assessments that involve human-AI collaboration. These rubric criteria are designed to help you assess the process of learning and evaluate whether students have appropriately developed and applied disciplinary skills and knowledge when they are working with AI.

Approaches to viva voces, live Q&A, and other interactive oral assessment

[Interactive oral assessments can be an authentic, secure, and engaging](#) way to assess attainment of learning outcomes. Optimally, they are conversational in nature, as opposed to a question and answer oral test. They allow you to probe deeper understanding – it is often very easy to spot a student who doesn't understand a concept by their oral responses. We have provided [guidance on these assessment tasks](#), and provide additional guidance below for interactive oral assessments in the context of AI.

Appendix 4 in the [full guide](#) has more guidance about these assessments.