

Generative Artificial Intelligence and Teaching and Learning at McMaster University

GENERATIVE ARTIFICIAL INTELLIGENCE AND TEACHING AND LEARNING AT MCMASTER UNIVERSITY

ASPENLED

CONTENTS

Introduction	1
--------------	---

Part I. Main Body

1. A Brief Introduction to Generative AI	3
2. Generative AI limitations and risks in post-secondary education	5
3. Generative AI within the Teaching and Learning Context of McMaster University	11
4. Designing Assessments in the Age of Generative AI	21
5. Generative AI and Student Learning and Perceptions	28
6. Using Generative AI as an Instructor	35

Generative artificial intelligence (generative AI) is changing how we teach and how we learn. What we want our students to learn – the core knowledge, skills and values of our disciplines – guide how we craft our curriculum and shape our pedagogical approaches. As educators we have long adapted what and how students learn to changing technology, changes in our disciplinary knowledge, and changes to the context of the University. We care about our students and about what and how they learn.

The capabilities of generative AI to produce coherent, logical and reflective text – as well as images, code, audio and video – invite new, and sudden, change to teaching and learning here at McMaster and around the world. How we respond to this change – if we respond – is a personal question, and an institutional one.

While many institutions and organizations are offering guidebooks, webinars and resources for adapting teaching methods and materials to address this rapid shift, the truth is we simply don't yet know the scale of change required. Will you want to adapt a single assessment? Will we need to rethink the core learning outcomes for a program? Will we need to reconsider the purpose of a post-secondary degree?

Media reports traverse the spectrum of panacea to catastrophe; conversations with colleagues and students here at McMaster mirror this breadth. Our individual reactions are shaped by our disciplinary backgrounds, our experience with generative AI and our teaching philosophies.

To say that any one guidebook – like this – can prepare you to teach amid the changes brought and coming by generative AI is foolish. We write this guidebook knowing some of its content will be obsolete in months. We wanted examples – so many examples – that we just do not have yet to offer (please: send us your examples!). We wanted to provide clear, simple and actionable advice for how to adjust your courses and your teaching methods, but ran up against the reality of idiosyncratic courses with unique assessments that each require slightly different guidance.

And so we offer this guidebook recognizing its limits. It aims to ground you in what generative AI is and what it might mean for student learning and for your teaching here at McMaster. It explores some of the ethical questions you may already be grappling with and invites you to share those we haven't yet considered. It offers specific advice for redesigning assessments and for how you might explore the use of generative AI in your teaching. And it tries wherever possible to be clear about what we don't yet know, but are trying to answer.

As authors we are educational developers, educators, researchers and students. We write this for you as colleagues and hope you will share with us your reactions, questions and suggestions. This guidebook will be updated – it will have to be updated – and we want to hear from you where we need to do more. You can reach us at mi@mcmaster.ca or mi.mcmaster.ca/aboutus

1.

A BRIEF INTRODUCTION TO GENERATIVE AI

Generative AI is a type of artificial intelligence that uses machine learning to generate new content by analyzing and processing vast amounts of data from diverse sources. Generative AI tools can generate text, images, video, sound and code. Different tools are trained on different datasets and with different training methods. The generated responses of these tools are probabilistic, which can result in errors in responses.

Large language models (LLMs), for instance, specialize in analyzing and processing text and generating new text. Different LLMs have distinct datasets and employ unique training methods. GPT 3 and GPT 4 are examples of LLMs. OpenAI's ChatGPT is a chatbot created on GPT 3 or GPT 4.

- A useful [glossary of AI terms can be found here](#) and a great brief introductory video from the Wharton School [can be found here](#)

While generative AI is not new, OpenAI's launch of ChatGPT in November 2022 [marked the fastest recorded adoption of a technology tool](#) to that date.

Over the intervening months, the release of similar text-based generative AI tools from Microsoft's Bing to Google's Bard, in addition to improvements in tools have contributed to a [perception of an explosion of AI](#).

Indeed, the rapid proliferation of tools and advancements in technology saw over 100 leaders in AI technology [write an open letter](#) urging a collective pause on AI developments more powerful than GPT 4 to give time for security and safety features to develop and for the creation of regulation and governance structures.

The need for such regulation or [governance extends to full nations](#), but also to specific sectors, such as post-secondary education, and in turn, McMaster University.

Broader issues related to generative AI include [privacy of personal data](#), [risks of misinformation](#), [existential risks](#), concerns about [job dislocation or loss](#), [environmental costs](#), [labour exploitation](#), and [copyright](#).

The next chapter in this book will look at some of these risks in more detail as they relate to post-secondary education.





An interactive H5P element has been excluded from this version of the text. You can view it online here:

<https://ecampusontario.pressbooks.pub/mcmasterteachgenerativeai/?p=186#h5p-1>

2.

GENERATIVE AI LIMITATIONS AND RISKS IN POST-SECONDARY EDUCATION

While the innovation and creativity of generative AI is exciting, these systems do not come without limitations or ethical challenges. Some of these challenges speak to the specifics of our post-secondary context – like academic integrity – while others intersect with communities, the environment, and humanity as a whole. Many AI experts have documented such alarming concerns including, [the size and scale of large language models, misinformation](#), [AI misalignment](#), and [existential risks to humanity](#).

For our purposes here this chapter will review some of the broader limitations and risks of generative AI, and will then turn to the specific context of post-secondary education. Many of these risks and limitations are explicitly addressed in [McMaster’s Provisional Guidelines on the Use of Generative AI](#), which will be addressed in more detail in the next chapter and should be regularly reviewed for updates.

General Limitations and Risks of Generative AI

Biases

Generative AI tools are trained on a range of data, some general models, like GPT-4 draw on a wide range of sources. Biases inherent in the training data – those that may discriminate against or marginalize underrepresented, minority, and equity-deserving groups – may appear in the results generated by these tools. While efforts have been made by companies like OpenAI to create ‘guardrails’ to prevent hateful and discriminatory results from being generated, the risk of bias persists in the limitations of the training data itself. That is to say existing biases in the training data may make a discriminatory result statistically more likely, and so the generative AI tool is more likely to produce that result. For instance, in a prompt to generate a story about slaying a dragon, the probabilistic result is to have a prince slay the dragon because that is the most common pattern in the training data. This somewhat innocuous example points to the broader risk of unexamined bias in generative AI results; that is to say, the result doesn’t have to be hate speech to be harmful, it does not have to be extreme to be biased.

As educators we need to be thoughtful about the ways these biases might be perpetuated or left unexplored when we use generative AI in teaching materials or student assessments. Teaching our students about the bias inherent in generative AI tools is one way we can collectively explore this limitation; we can also invite students

to reflect on and discuss biases they notice in generative AI outputs and to comment on how and why these biases may appear.

Hallucinations

Generative AI tools make things up. As probabilistic models they are designed to generate the most likely response to any given prompt. Given that these tools do not ‘know’ anything, and are – in most instances – limited in their ability to fact check, the responses they generate can include factual errors and invented citations/references. This known phenomenon has been termed ‘hallucination,’ and is one persuasive reason to evaluate and fact-check all responses a generative AI tool produces.

There is some speculation that generative AI tools will soon include a ‘confidence indicator’ that might let users know the degree of confidence the tool has that a generated response is accurate. Likewise, some reporting suggests that generative AI tools will begin to fact-check their responses against internet sources or other AI models. At the time we are writing, these capabilities are not in wide circulation. Instead, we need to teach our students – and practice ourselves – healthy skepticism about the reliability of generative AI produced responses and a consistent practice of checking outputs against verified sources.

Environmental and Human Costs

The exact environmental costs of generative AI models is hard to know. The size of the model, the training approach used and the capabilities of the tool influence how much energy the model uses. Likewise, there are very different energy needs for training a model and for using it. Some prominent companies deploying generative AI tools – like Google and Microsoft – have also pledged to be carbon neutral or carbon negative in a way that – ostensibly – accounts for the energy use of their generative AI tools Bard and Bing, respectively. That said, the known energy-consumption of these tools should not be a limitation left unexamined. The energy costs of training and running the tools – while not definitively known – [is estimated to be considerable](#).

As educators we have a responsibility to share this impact with our students; as a community at McMaster we have an opportunity to make a difference by contributing to carbon offsetting programs and to educating our students on the environmental cost of these tools.

Just as there is variation in the environmental impact of generative AI tools based on their size and capabilities, there is variation in how these models are trained. Some tools, like ChatGPT, have been trained using ‘reinforcement learning through human feedback.’ This kind of training for the model involves humans reviewing a prompt and the generated output and ranking or ‘up or down voting’ in a way that gives the model feedback about the accuracy and helpfulness of the generated output. In addition to training the accuracy of outputs, workers are also used to review outputs against guardrails of appropriate content or “content moderation.” While technology tools, including social media and generative AI, have long employed human workers for content moderation, OpenAI came under criticism for outsourcing this practice to [low-wage](#)

[workers in Kenya](#). These workers must sift through toxic and explicit content with an aim of creating safer systems for the broader public without full consideration of psychological wellbeing.

Disinformation

The ability of generative AI to create realistic and plausible text, video, audio and code makes the creation of false, biased, or politically motivated media faster and easier to produce. Our individual and collective ability to identify reliable and trustworthy sources, and to evaluate what we read, view and hear has never been more important.

Copyright

As mentioned, generative AI tools use a wide range of data to learn from before producing outputs. Many of these tools include in their datasets content created and shared publicly – like Twitter or Reddit – as well as that created by artists or users without explicit consent for inclusion in a dataset for generative AI use. Ongoing lawsuits related to copyright filed by artists are challenging the inclusion of creative works in datasets, and [open questions remain about what might be fair use](#). Some tools, like Adobe Firefly, are working to compensate contributors, while others are tailoring their datasets to include content with consent for inclusion explicitly obtained.

Privacy and Safety

Without consistent government regulation of emerging generative AI tools, users rely on the user agreements and privacy guidelines of specific tools. Here at McMaster we have privacy and security protocols that see technology tools routinely evaluated for privacy and security risks. At the time of writing a complete privacy and security assessment of generative AI tools has not been completed. As such, we recommend that users – including educators and students – carefully review user agreements and understand the ways in which generative AI tools may collect and make use of user data before consenting to use of the tools. Such transparent communication and shared discussion can both educate students on these agreements, and potentially alleviate concerns about privacy and safety.

Of particular note in our context is the limitation on user age for ChatGPT – users must be 18+ to create a user profile.

Many generative AI tools, including ChatGPT, have settings that allow for users to turn off data collection, which means the tool will not use the inputted prompts or data for later use.

Reading and reviewing user and privacy agreements yourself and with your students is one concrete way to understand the implications of these tools on data privacy and security. You can also opt to [turn off data collection](#) on those tools that permit it, and share this practice with your students.

Post-secondary Specific Limitations and Risks

While the limitations and risks outlined earlier in this chapter also apply to the post-secondary context, there are several risks specific to our University environment worth considering, specifically supporting academic integrity and equitable access.

(Re)defining academic integrity and academic dishonesty

McMaster's [academic integrity policy](#) defines academic dishonesty as “to knowingly act or fail to act in a way that results or could result in unearned academic credit or advantage” and that “it shall be an offence knowingly to ... submit academic work for assessment that was purchased or acquired from another source.”

In an article describing how he integrated generative AI into writing assignments, Paul Fyfe observes “computer- and AI-assisted writing is already deeply embedded into practices that students already use. The question is, where should the lines be drawn, given the array of assistive digital writing technologies that many people now employ unquestioningly, including spellcheck, autocorrect, autocomple, grammar suggestions, smart compose, and others [...] within the spectrum of these practices, what are the ethical thresholds? At what point, in what contexts, or with what technologies do we cross into cheating?” He continues, “educational institutions continue to define plagiarism in ways that idealize originality” ([Fyfe, 2022](#)).

In this observation, Fyfe highlights a recurring theme in the literature around academic integrity and artificial intelligence, that is: with these technologies the defined boundaries of independent work have become porous, and the contrast between “humanity originality and machine imitation” (Fyfe, 2022) blurs.

The result of this shift in understanding is a call within the literature to reexamine, and perhaps redefine, what constitutes plagiarism, academic integrity and academic dishonesty, with some authors arguing that “Academic integrity is about being honest about the way you did your work” ([David Rettinger in Surovell, 2023](#)), others urging a defended boundary of primarily individual effort ([Keegin, 2023](#)), and still others arguing for a new framework entirely – what [Sarah Eaton](#) calls ‘post plagiarism’ through a norm of human hybrid writing.

Where most of the reviewed literature holds consensus is that using generative artificial intelligence does not automatically constitute academic misconduct ([Eaton, 2022](#)), but rather, to quote the European Network for Academic Integrity, “Authorised and declared usage of AI tools is usually acceptable. However, in an educational context, undeclared and/or unauthorised usage of AI tools to produce work for academic credit or progression (e.g. students’ assignments, theses or dissertations) may be considered a form of academic misconduct” ([Foltynek, 2023](#)).

Detection

Questions around detecting AI generated writing fall into (1) the technological – is it possible to reliably detect AI-generated writing? (2) the philosophical – is the role of the educator one of trust or one of surveillance? and (3) the existential – what is the value of a university degree if the academic labour behind it is uncertain?

There are not yet reliable detection tools. Those that are available – GPTZero, Turnitin, Originality.ai, etc – have been found to [misidentify original student content](#) as AI generated, with some findings demonstrating that “these detectors consistently misclassify non-native English writing sample as AI-generated, whereas native writing samples are accurately identified.” (Liang, 2023). Moreover, students have not consented to having their work submitted to these tools, with open questions related to data privacy and security (Mortati, 2021).

While technology and a perceived ‘arms race’ between detection and AI tools pose their own challenges (Mortati, 2021), there are also questions about the role of educators and their assumptions about students as learners. With significant evidence pointing to student academic misconduct on the rise, [particularly over the pandemic](#), there are arguments that “we must prioritize student learning above catching cheaters” (Eaton, 2022) and that understanding why students engage in academic misconduct may point to approaches to reduce these behaviours. Indeed, the instances of academic dishonesty and opportunities to cheat predate generative AI; what the tools introduce is “ease and scope” (Supiano, 2023) that amplifies an existing challenge.

[Students’](#) self-reported reasons for academic misconduct include performance pressure, high stakes exams, overwhelming workload, being unprepared, feeling ‘anonymous’, increased opportunities to cheat enabled by technology, peer acceptance of cheating, misunderstanding plagiarism, and feeling like it will go unpunished. This [research brief](#) on why students cheat summarizes research findings that argue for a reduction in academic dishonesty when students are both clear about what constitutes academic integrity/academic dishonesty, what the expectations are for their academic work and a felt perception of mutual benefit for behaving with integrity rather than competition with other students. In short [“Students are more likely to engage in academic misconduct when they are under pressure, when there is an opportunity, and when they are able to rationalize it.”](#)

Instead of positioning the educator as one to detect and survey, these pieces suggest the role be one of designing authentic and scaffolded assessments and explaining and exploring academic integrity with students.

Within these proactive strategies for cultivating academic integrity is an implied sense of time and scale – that is, these strategies imagine instructors have sufficient time, resources and energy to update or redevelop courses and assessments. [Providing scalable, supported and realistic assessment redesign](#) will be one of the ongoing areas of need for educators as generative AI is integrated into more tools and more courses. A later chapter

in this book focuses specifically on strategies you might take to redesign assessment to promote academic integrity.

Equitable Access

Cost of tools poses a barrier for many students in accessing generative AI tools. With many tools currently available for free, some of these – like ChatGPT – have paid tiers with significant improvements in functionality and performance for paid subscribers. Those students who can afford to pay for paid tiers may be disproportionately advantaged in assignments that incorporate the use of generative AI. As educators we need to design activities that encourage the use of free versions. For instance, Microsoft’s Bing, used in creative mode, draws on GPT-4, the same model that powers the paid version of ChatGPT. Designing assessments that draw on these free versions will make access for all students easier, even while there are continuing inequities in terms of internet availability, cost and speed.

That said, if students are learning online from other countries, some particular tools, like ChatGPT may be restricted due to government regulation or censorship. Attention to this possibility may mean allowing some students to opt-out of assignments that use generative AI, or providing alternatives for their engagement.

Finally, the intersection of generative AI and students with disabilities is an area of emerging research; we aim to add more information about generative AI as assistive technology in the coming months.

3.

GENERATIVE AI WITHIN THE TEACHING AND LEARNING CONTEXT OF MCMASTER UNIVERSITY

As the previous chapter highlighted, the risks and challenges to post-secondary education created or energized by generative AI are significant and wide-reaching from how we assess student learning, to promoting academic integrity, to considering what we want students to learn and what their future will be after graduation. Taken together these challenges are significant, and required a full institutional response. This chapter reviews how McMaster responded to this challenge, and what will come next as an institution.

Generative AI at McMaster University

While generative AI emerged as a transformative technology tool after McMaster's Institutional Priorities and Strategic Framework (2021-2024) and McMaster's Partnered in Teaching and Learning Strategy (2021-2026) were launched, the impact of generative AI nevertheless aligns with existing institutional strategic priorities and ongoing efforts to enhance teaching and learning.

In [McMaster's Institutional Priorities and Strategic Framework \(2021-2024\)](#), for example, Teaching and Learning, one of five priorities listed, identifies the development of active and flexible learning spaces as one key objective. It notes that in "recognizing the ways that online and virtual classrooms have changed the teaching and learning environment for both our educators and our students, [McMaster must] use evidenced based research to make decisions about tools and platforms to optimize learning in the digital environment" (p. 10). While there is little peer reviewed literature yet available on generative AI in post-secondary teaching and learning, McMaster is staying abreast of such research, and even engaging in research of its own in an effort to develop and maintain guidelines and good practices with respect to the usage of generative AI at the institution.

Likewise, [McMaster's Partnered in Teaching and Learning Strategy \(2021-2026\)](#) connects to generative AI via not one, but two of its four strategic pillars: 1) Fostering Inclusive Excellence and Scholarly Teaching strategy, via the themes Teaching as a Professional and Innovative Practice, and Assessment and Evaluation of Student Learning, and 2) Developing Active and Flexible Learning Spaces, via the Digital Learning theme.

McMaster's Task Force on Generative AI in Teaching and

Learning

Recognizing that the initiatives in these strategies alone could not respond quickly enough to the challenges presented by generative AI, on May 1, 2023, McMaster University struck a Task Force on Generative AI in Teaching and Learning to consider impacts posed by generative AI on teaching and learning at McMaster. The Task Force was also to provide strategic guidance and actionable recommendations for educators planning for fall courses. Co-chaired by Kim Dej, Vice-Provost, Teaching and Learning, and Matheus Grasselli, Deputy Provost, the Task Force includes students, faculty, and staff from across the university. Recommendations made by the Task Force will be submitted to Susan Tighe, Provost and Vice-President (Academic) in the fall of 2023.

The following overarching principles have guided the work of the Task Force and will continue to be updated through conversations with the McMaster campus community.

- Students want to learn, and instructors want to support their learning.
- Participatory learning – learning which happens in relationships and community – continues to be a valuable and vital way for students to learn.
- Assessments that require students to document the process of learning continue to be meaningful for student learning.
- Generative AI poses risks, as well as opportunities. Individuals will have different reactions and different expectations for the technology.
- Disciplinary differences and departmental cultures will vary around the use of generative AI.

On June 5 the Task Force released [Provisional Guidelines: The Use of Generative Artificial Intelligence \(AI\) in Teaching and Learning at McMaster University \(June, 2023\)](#) for McMaster students and educators. The guidelines are intended to offer a starting point for instructors to understand the potential uses of generative AI in their teaching and student learning and for developing courses for the fall term.

These guidelines will continue to be updated as the Task Force explores additional topics and as technology changes. Members of the Task Force invite feedback and suggestions on these guidelines through [this form](#). It is expected these guidelines will be updated again in time for winter course preparation. Potential policy changes implied by these guidelines will be addressed by the relevant governance bodies.

Staff at the MacPherson Institute are available to consult with instructors regarding these guidelines; Instructors can email mi@mcmaster.ca for support.

Provisional Guidelines: The Use of Generative AI in Teaching and Learning at McMaster

1. Instructors are not required to use generative AI tools for teaching.
2. McMaster's existing academic integrity policy applies when using generative AI. Its overall definition of academic dishonesty, which is to knowingly act or fail to act in a way that results or could result in unearned academic credit or advantage, allows for allegations related to generative AI. The policy states under item 18(c) that "It shall be an offence knowingly to ... submit academic work for assessment that was purchased or acquired from another source".
 - a. Unless otherwise stated, students should assume use of generative AI is prohibited.
 - b. Instructors who incorporate generative AI into courses should explain to students in writing and verbally in-class how generative AI material should be acknowledged or cited (see Appendix A for examples).
 - c. Updated guidance on instructor use of generative AI for feedback and grading will be provided by the fall.
3. Individual instructors should determine if generative AI will be incorporated into course design, activities, and assessments based on course learning outcomes, individual interest, and conventions and expectations of the discipline.
4. Individual instructors should clearly communicate to students if and to what extent generative AI is acceptable in the course in the course outline, verbally in-class and in assessment descriptions (see Appendix B for examples).
5. If instructors use generative AI in their teaching materials instructors should explain in the course outline the extent to which generative AI has been, or will be, used.
 - a. Instructors should fact-check any generative AI produced materials.
 - b. Instructors should not submit student work to generative AI tools for feedback without students' consent and ability to opt-out.
6. Instructors incorporating generative AI should be aware of the privacy policies and user agreements of each generative AI tool and alert students to these policies in the course outline.
7. Where possible, courses that incorporate generative AI should rely on free versions of generative AI tools (e.g. Microsoft Bing, ChatGPT 3) for student use.
 - a. Alternatives should be provided for Generative AI tools that are restricted to users 18+ (e.g. ChatGPT).
8. Generative AI plagiarism detection software is currently unavailable or not recommended at McMaster. This software will continue to be reviewed and may be used in the future.
 - a. These detectors will produce false positives and are not approved for use through the University's

policy. Students have not consented to the sharing of their intellectual work through these tools. It is also unclear how the material submitted to the third-party detectors is retained or used.

b. Until more is understood about generative AI detection tools, instructors should not submit student work to generative AI detection tools.

c. McMaster has an institutional membership to Turnitin, a plagiarism detection software. Turnitin announced an update aimed at detecting writing produced by generative AI. McMaster, like many other institutions, has not yet turned on this feature as there is a need to understand the functionality of the tool, assess the security and privacy considerations for student work and determine whether it aligns with existing policies. **d.** If you do suspect student work may have violated the academic integrity policy, please review the steps to take.

9. Instructors with courses that incorporate generative AI should:

a. Consider the course learning outcomes and ensure the incorporation of generative AI will support core learning outcomes; and ensure incorporation offers meaningful learning, rather than inclusion for the sake of novelty.

b. Describe or discuss with students the strengths, limitations and ethical considerations of the technology, including factual inaccuracies or ‘hallucinations’, societal biases present in the training data and the rationale for using generative AI in assignments. (see Appendix A for examples).

c. Resources for faculty to talk with students about generative AI are in development and will be available by late August.

10. Instructors who include assessments that incorporate generative AI should:

a. Consider including reflective components that invite students to comment on the use of/experience with generative AI in the assessment

b. Explicitly review criteria and/or rubrics in ways that demonstrate how the use of generative AI is being assessed (see Appendix C for example).

11. Assessment alternatives that may be less susceptible to the use of generative AI include oral exams, presentations followed by a Q and A, invigilated/in-class assessments, practical tests, assessments that incorporate class discussion/activities, and process-based work.

12. Instructors may consider adding an honour pledge (see Appendix D for example) to assessments.

13. The MacPherson Institute will continue to provide training and resources for instructors and students on how to use generative AI effectively. See mi.mcmaster.ca for current workshops, resources and to schedule a consultation.

14. McMaster will explore an annual donation to carbon offsetting programs to address the environmental impact of training large AI models.

15. The MacPherson Institute will collect feedback from instructors and students this fall on their experiences, questions and concerns about using generative AI in teaching and learning in an effort to update and improve these guidelines.

16. These guidelines will be regularly reviewed and revised with the aim of updating them before winter course outlines are due.

17. Course instructors have three options for directing teaching assistant use of generative AI:

a. Permitting teaching assistants to use generative AI for any aspect of teaching assistant work, with the exception of summative evaluation, with no expectation that they use generative AI and no training specific to generative AI required. TAs must inform the instructor of the intended use of generative AI, and receive approval, before implementation. Summative evaluations are those which significantly impact a student's grade or progress in a course. This includes providing a quantitative grade (number or letter grade).

b. Requiring teaching assistants to use generative AI for specified teaching tasks as outlined in the hours of work form and with training provided. In the instance of required use: As directed by the course instructor explicitly in the hours of work form, teaching assistants will use generative AI for the specific teaching tasks. Course instructors will provide teaching assistants with the necessary training to use generative AI for the specified teaching purpose(s) with this training included in the hours of work. Teaching assistants will evaluate all teaching materials/formative feedback developed with generative AI for accuracy before use with students. Any planned use of generative AI by teaching assistants will be shared with students in the course outline.

c. Prohibiting teaching assistants from using generative AI for teaching tasks

18. Generative AI tools can be used to provide formative feedback on student work; generative AI tools cannot be used to provide summative evaluation of student work.

- AI-generated formative feedback is intended to guide learning and improve understanding, by pointing out strengths and areas for improvement in student work.
- Summative evaluations are those which significantly impact a student's grade or progress in a course. This includes providing a quantitative grade (number or letter grade).

19. Instructors, or teaching assistants when directed, should review AI-generated formative feedback to ensure it aligns with the learning objectives and course materials, and add their own insights where necessary. Formative feedback that uses AI should not be given a quantitative grade by the AI tool. A "pass/fail" or "completion" may be applied.

20. Instructors, or teaching assistants when directed, are responsible for summative evaluations to ensure appropriateness and accuracy. Data collection should be turned off on generative AI tools when used for providing formative feedback. Ongoing work to complete privacy impact assessment and security evaluation on recommended generative AI tools will be communicated with the campus community.

21. When providing AI-generated formative feedback, students should be made aware that it is generated by AI explicitly in the course syllabus.

22. Students may opt-out of assessments that require the use of generative AI only in exceptional circumstances as approved by the course instructor. If approved to opt-out of an assessment that requires the use of generative AI based on an exceptional circumstance, students will not face academic penalty, but will be required to provide alternative and equivalent evidence of their learning as proposed to, and agreed to by, the course instructor.

Appendix A: Citation and Reference Guidelines

A McMaster specific citation guide is in development through the Library. Until then, please consider citation options such as:

“[Generative AI tool]. (YYYY/MM/DD of prompt). “Text of prompt”. Generated using [Name of Tool.] Website of tool”

E.g. “ChatGPT4. (2023/05/31). “Suggest a cookie recipe that combines oatmeal, chocolates chips, eggs and sugar.” Generated using OpenAI’s ChatGPT. <https://chat.openai.com>”

Instructors may also consider requiring students to include a reflective summary at the end of each assessment that documents what generative AI tools were used, what prompts were used – including a complete chat log – and how generated content was evaluated and incorporated.

Other citation guidelines can be viewed at:

- [MLA Guidelines](#) on citing generative AI
- [APA Guidelines](#) on citing generative AI
- [Chicago FAQ](#) on generative AI
- A [quick guide provided from the University of Waterloo](#), with a McMaster version coming in Fall 2023.

Appendix B: Sample McMaster Syllabus Statements

Use Prohibited

Students are not permitted to use generative AI in this course. In alignment with [McMaster academic integrity policy](#), it “shall be an offence knowingly to ... submit academic work for assessment that was

purchased or acquired from another source”. This includes work created by generative AI tools. Also state in the policy is the following, “Contract Cheating is the act of “outsourcing of student work to third parties” (Lancaster & Clarke, 2016, p. 639) with or without payment.” Using Generative AI tools is a form of contract cheating. Charges of academic dishonesty will be brought forward to the Office of Academic Integrity.

Some Use Permitted

Example One

Students may use generative AI in this course in accordance with the guidelines outlined for each assessment, and so long as the use of generative AI is referenced and cited following citation instructions given in the syllabus. Use of generative AI outside assessment guidelines or without citation will constitute academic dishonesty. It is the student’s responsibility to be clear on the limitations for use for each assessment and to be clear on the expectations for citation and reference and to do so appropriately.

Example Two

Students may use generative AI for [editing/translating/outlining/brainstorming/revising/etc] their work throughout the course so long as the use of generative AI is referenced and cited following citation instructions given in the syllabus. Use of generative AI outside the stated use of [editing/translating/outlining/brainstorming/revising/etc] without citation will constitute academic dishonesty. It is the student’s responsibility to be clear on the limitations for use and to be clear on the expectations for citation and reference and to do so appropriately.

Example Three

Students may freely use generative AI in this course so long as the use of generative AI is referenced and cited following citation instructions given in the syllabus. Use of generative AI outside assessment guidelines or without citation will constitute academic dishonesty. It is the student’s responsibility to be clear on the expectations for citation and reference and to do so appropriately.

Unrestricted Use

Students may use generative AI throughout this course in whatever way enhances their learning; no special documentation or citation is required.

Appendix C: Sample Rubrics

Sample Rubrics Developed with ChatGPT:

I acknowledge the use of ChatGPT 4.0 to create sample analytic and holistic rubrics. The prompts included “Imagine you are a rubric generating robot who creates reliable and valid rubrics to assess university-level critical thinking skills. You have been tasked with generating a rubric that evaluates students critical thinking skills and incorporates their use of generative AI. Create two holistic rubrics and two analytic rubrics to assess

these skills.” The output from these prompts was to provide examples of the kind of rubrics that could be used to assess the integration of generative AI in course assignments.

Rubric 1: Assessing Generative AI Use and Integration

Criteria	4	3	2	1
Argument Structure	The argument is clearly articulated and logically structured.	The argument is generally clear and logical, with minor inconsistencies.	The argument is somewhat unclear or inconsistently structured.	The argument lacks clarity and logical structure.
Evidence	Evidence is thorough, relevant, and convincingly supports the argument.	Evidence is generally strong and relevant, with minor lapses.	Evidence is somewhat sparse, irrelevant, or does not fully support the argument.	Evidence is lacking or largely irrelevant.
Use of Generative AI	AI is used effectively to support arguments, demonstrating a high understanding of its capabilities and limitations.	AI is used effectively, but understanding or integration could be improved.	AI is used, but not effectively integrated or misunderstood.	AI is not used or its use does not contribute to the argument.
Reflection on AI	The student clearly articulates how AI contributed to their critical thinking process and considers its limitations.	The student generally explains how AI contributed to their thinking, with minor lapses in considering its limitations.	The student’s explanation of how AI contributed to their thinking is unclear or superficial.	The student does not explain how AI contributed to their thinking.

Rubric 2: Assessing Generative AI Use and Integration

Criteria	4	3	2	1
Understanding of AI	The student demonstrates a deep understanding of the capabilities and limitations of the AI.	The student demonstrates a good understanding of the AI, with minor misconceptions.	The student shows a basic understanding of the AI, but has significant misconceptions.	The student shows little to no understanding of the AI.
Integration of AI	AI is seamlessly integrated into the work, effectively augmenting the student's critical thinking.	AI is generally well integrated, though at times it may seem somewhat forced or awkward.	AI integration is inconsistent or superficial, not effectively augmenting the critical thinking process.	AI is not effectively integrated into the work.
Reflection on AI	The student clearly reflects on the role of AI in their work, considering both its contributions and its limitations.	The student generally reflects well on the AI's role, though considerations of its limitations may be superficial.	The student's reflection on the AI's role is minimal or lacks depth.	The student does not reflect on the AI's role in their work.
Innovation with AI	The student uses AI in novel or innovative ways to enhance their argument.	The student uses AI effectively, though it may lack innovation.	The student uses AI in a straightforward or predictable way, not enhancing the argument.	The student does not use AI in an innovative or meaningful way.

Appendix D: Honour Pledges

Honour pledges are formal, student-led commitments to uphold the principles of academic honesty and integrity. These pledges represent students' personal assurance to maintain and respect academic standards, abstaining from any form of plagiarism, cheating, or other academic misconduct. They often form part of the assessment submission process, where students attach a pre-defined pledge to their work as a statement of authenticity. [Several studies](#) have investigated the use of honour codes and academic integrity and found them effective in reducing academic dishonesty.

Instructors might consider developing honour pledges together with their students, or adapting this McMaster honour pledge to their purposes.

"I understand and believe the main purpose of McMaster and of a university to be the pursuit of knowledge and scholarship. This pursuit requires my academic integrity; I do not take credit that I have not earned. I believe that academic dishonesty, in whatever form, is ultimately destructive to the values of McMaster, and unfair to those students who pursue their studies honestly. I pledge that I completed this assessment following the guidelines of McMaster's academic integrity policy."

Forthcoming Guidelines and Resources

Staff across the University are working to develop additional resources and supports.

Some known needs from the campus community include:

- Resources for students to understand generative AI risks and opportunities
- Overview of generative AI tools including privacy and security assessments
- Digital literacy learning outcomes and digital literacy resources

4.

DESIGNING ASSESSMENTS IN THE AGE OF GENERATIVE AI

Chapter Contents

[Key Assessment Design Principles](#)

[Part 1: Shorter-Term Approaches](#)

[Part 2: Redesigning Your Assessment](#)

Introduction

Assessment has become a thorn in the side of many educators over the past three years. First, the rapid shift to remote teaching during the pandemic forced many educators to adopt assessment approaches that they may not have been comfortable with or that they recognized were not ideal for student learning. Then – just as many of us were returning to the more familiar assessment circumstances of in-person classes – OpenAI released ChatGPT. Any assessment with a non-invigilated written component, including the writing of computer code, now raises questions about if – and to what extent – students are making use of generative AI.

At the same time, the events of the past three years have highlighted existing troubles with our assessment practices and prompted us to reflect on the purpose of assessment in teaching. The language of care in teaching that became more prevalent during the pandemic helped to reframe the conversation about academic integrity into a deeper consideration of why students cheat. One culprit is poorly designed assessments, which may:

- only require students to recall what they have already learned, and/or
- are mismatched with what students expect to do and learn in the course, and/or
- unfairly disadvantage some students and not others, and/or
- have unnecessarily high stakes.

Assessment (re)design thus offers educators the opportunity to have a meaningful impact on issues of academic integrity.

We recognize that you may not currently be able to significantly redesign your course assessments, which takes time and effort to do thoughtfully. Trying anything new in the classroom also carries a degree of risk, particularly for educators who are already in precarious roles like sessional instructors, contractually limited appointments and pre-tenure faculty. Even if you do have the capacity to redesign your assessments, we suggest starting small: addressing the assessment that concerns you the most or will have the greatest impact, and then building on your experience.

We have divided the chapter into two parts:

1. a series of shorter-term, “quick fix” strategies to help counteract or embrace the easy access to generative AI, and
2. a workbook to guide you through the redesign of an assessment, based on our intensive Assessment Development Workshop.

We hope there is a path through the resource for all educators, acknowledging that you will each be teaching in different contexts, be at different points of your career, and be working under different conditions. We also encourage you to not go through the resource in isolation but rather, to reach out to your Faculty’s key contact at the MacPherson Institute and discuss your assessment further.

Key assessment design principles

Throughout the chapter, we will foreground the following assessment design principles:

- Authentic assessments
 - Assessments are authentic when they “replicate real world performances as closely as possible” (Sviniki 2004), “foster[ing] disciplinary behaviours and ways of thinking and problem solving used by professionals in the field” (via Queen’s U module); we will elaborate on what defines “authentic” assessments later in the chapter.
- Learner-centeredness
 - Assessments reflect the goals, interests and lived experiences of the learners;

learners can see themselves in the assessment and are intrinsically motivated to complete it well.

- Universal design for learning (UDL)
 - Assessments are proactively designed with accessibility in mind, with the aim of eliminating barriers to give all students an equal opportunity to succeed.
- Constructive alignment
 - Assessments are aligned with course and program learning outcomes; i.e., does the assessment demonstrate that the learner has met the course learning outcomes? Do learners evidence their grasp of essential course skills and knowledge by completing the assessment?
- Assessment *for* learning
 - Assessments are opportunities for students to enrich and extend what they have learned by applying them in novel contexts; the assessment itself is a site of learning.

Part 1: Shorter-term approaches

If redesigning your assessments is not feasible for you, you might want to consider shorter-term strategies to tweak existing assessments. It's important to recognize that these are 'short-term' and that they may not be as effective or sustainable as AI capabilities improve (and they are improving at a fast pace!). We've included the following options that are more easily integrated into your existing assessments:

- Invigilated / Observed In-class work
- In-class work that integrates AI
- Revising assessments to emphasize tasks AI cannot perform well
- Revising grading schemes and rubrics

Be mindful of how an in-class assessment might present barriers for students – particularly those that may have academic accommodations in place. A well-structured, timed, writing exercise, for example, may cause significant concern for students who struggle with cognitive load or focus issues ([Reference](#)).

Invigilated/observed in-class work

Traditionally, postsecondary has leaned heavily on invigilated tests and exams. A well-intended effort to move away from these to more authentic assessments has required much thought and labour on the part of the instructor, who perhaps now are grappling with how it may have unintentionally created new ways for students to use AI to complete the assignment. One option to circumvent a possible plagiarism risk is to introduce space during in-class time for assessment of learning. This may involve shifting some of the content delivery to an asynchronous environment (e.g., recorded lectures, assigned readings). Some ideas may include:

- Real-time in-class discussions and reflections – an in-class group discussion, individual written reflection or oral response to a topic-based prompt demonstrates critical thinking and personal reflection.
- Group case studies – tapping into the proven practice of Problem-Based Learning (PBL), present students with real-life scenarios or case studies and use in-class time to collectively discuss, apply knowledge and problem-solve to analyze, propose solutions and back-up choices. Even as an ‘out of class’ activity, this approach requires the human judgement and contextual understanding that makes it less susceptible to AI-shortcuts.
- Presentations and debates – assigning topics or issues to research in class and present in the form of short presentations or paired/group debates can provide an opportunity to assess information literacy, argument structure and persuasive communication skills in addition to knowledge of the topic in a way that circumvents tasks that might be more easily automated by AI. Including a Q and A as part of a presentation also invites an opportunity for dialogue and engagement on what the student has learned.

In-class work that integrates AI

Integrating AI into low-stakes in class assessments will help communicate to your students that you recognize how AI can be used, and at the same time foster a better understanding for your students around how it can be used intentionally, ethically, and in support of your teaching goals. Here are just a few ideas for in-class activities:

- Instead of starting the class with a question on a key course topic, include the AI-generated answer and invite the class to critique and revise – either independently or in small groups.
- Hold a ‘humans vs AI’ debate where students pit their answers to topic prompts against those generated by AI. This discussion will help students organize their arguments’ points of view and also discover new perspectives and strengthen critical analysis skills.
- Divide the class into three groups and have each group evaluate the AI-generated output to a course topic prompt for either [factual accuracy, artificial empathy or bias](#).

Emphasizing tasks that AI cannot perform well

In the same vein that educators try to design authentic assessments that are valuable learning opportunities with a side benefit of being not easily plagiarized through the affordances of the pre-AI Internet, we have focused on trying to come up with alternative assessments that ChatGPT cannot easily perform. This is risky, as the technology is evolving at a quick pace, with massive amounts of prompts being continually added, and thereby improving the outputs being generated. However, while AI-proofing assessments may be nearly impossible, we can choose to focus on tasks in assessment expectations that encourage personalized and localized connections. An example of this might be to demonstrate the learning through links to local context, current events which may not be well represented in LLMs.

Revising grading schemes and rubrics

You may want to revise your current grading scheme, rubrics and the criteria outlined to reweight and emphasize the less-mechanical (aka easily AI created) competencies. For example, with a writing assignment you may choose to focus more on skills such as creating a good argument, including and evaluating evidence and critical analysis rather than grammar, and essay structure. You may also require that rough planning notes be included as part of their assignment submission to indicate assessment of the process rather than the finished product.

Part 2: Redesigning Your Assessment

If you have the time and capacity to redesign an assessment or two, read on! What follows is a set of considerations in the form of a workbook. Its purpose is to guide you through developing an engaging, meaningful and relevant assessment that students see value in completing – without the inappropriate use of generative AI – to further their learning, experience and expertise.

How to use the workbook

The workbook contains a series of reflective prompts, with fillable form fields to record your thoughts that you can then use in the design or redesign of your assessment.

You can choose to fill out the workbook prompts from start to finish, which will involve a considerable time commitment, or you can equally opt to focus on the prompts that will be the most helpful to your design process – though we do encourage you to read through the various considerations to determine whether they apply to your assessment.

When you have filled out the applicable prompts in each section of the workbook, you can use the “Export workbook responses” button to output them as an MS Word Document. To allow for the maximum flexibility of use, the button only outputs responses for the current section of the workbook – if you would like to keep all of your responses together, you can copy-and-paste them into the same document.

Getting started



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://ecampusontario.pressbooks.pub/mcmasterteachgenerativeai/?p=34>

Remember to export your workbook responses before moving on if you would like to keep them for your later reference!

Authentic assessment



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://ecampusontario.pressbooks.pub/mcmasterteachgenerativeai/?p=34>

Remember to export your workbook responses before moving on if you would like to keep them for your later reference!

Other assessment design considerations



One or more interactive elements has been excluded from this version of the text. You can view them online here: <https://ecampusontario.pressbooks.pub/mcmasterteachgenerativeai/?p=34>

Remember to export your workbook responses before moving on if you would like to keep them for your later reference!

Next Steps

With your workbook complete to the extent you wish, we encourage you to use it as a departure point for redesigning your assessment. The next steps are up to you!

You may wish to take the workbook to a consultation with your Faculty's [key contact](#) at the MacPherson Institute, or as tool for individual reflection and development. The Assessment Development Workshop, on which the workbook is based, involves peer feedback – consider finding a colleague who is similarly interested in redesigning an assessment. You can support each other through the process!

5.

GENERATIVE AI AND STUDENT LEARNING AND PERCEPTIONS

Generative AI and Opportunities for Student Learning

As McMaster has taken the approach of having each instructor decide whether and how to incorporate generative AI into a course or assignment, you may be wondering why you might want to do so. What benefits, if any, does generative AI pose for student learning? What learning outcomes could its use support or enhance? This chapter assumes your familiarity with the risks and challenges of generative AI for post-secondary (e.g. academic integrity, assessment design, hallucinations) and imagines what benefits their might be and what opportunities for preparing students for a generative AI supported learning experience.

You can think of the possibilities in two domains: (1) supporting personalized learning and (2) generating academic content.

Generative AI has many capabilities in supporting personalized learning, some of which we detail below. Chief among them is providing actionable, timely and relevant feedback on drafted student content. This feedback might be focused on the grammar or style of the draft, or on the logic of the argument, organization of the piece, or further examples to consider.

With respect to generating academic content or performing academic skills, you want to think carefully about what the core learning outcomes are for the course, and whether and how students can demonstrate these outcomes. Those skills or knowledge that are not essential to the core learning outcomes might be appropriate for ‘cognitive offloading’ to a generative AI tool. Cognitive offloading refers to the use of external resources or tools to change the information processing requirements of a task so as to reduce cognitive demand ([Risko and Gilbert, 2016](#)). For instance, if your course learning outcomes require students to demonstrate abilities to generate multiple hypotheses to explain a phenomenon, using generative AI to generate these hypotheses would be inappropriate. However, if your course learning outcomes were focused on having students test a hypothesis in a laboratory setting, having a generative AI tool generate the hypothesis which the student would then test would be an example of appropriate cognitive offloading.

In what follows we offer some concrete examples of how generative AI can be used to support personalized learning or generate academic content. For each example we remind you of the importance of first deciding whether this is a task the student needs to complete themselves to fulfill the course learning outcomes, or whether this is a task that would benefit from cognitive offloading. If you have questions or want to discuss, please reach out to an educational developer at the MacPherson Institute at mi@mcmaster.ca

Supporting Personalized Learning

Invite students to use a generative AI tool to:

- **check grammar and improve paragraphing** in order to detect and fix errors. This capability may be particularly helpful for non-native speakers, who may compose technically correct but awkward phrasing or use unconventional grammar structures. Students can also ask the generative AI tool to explain why a grammar mistake happened or why a sentence has been rephrased; this explanation can help students learn the relevant grammar rule and improve their overall writing. [Example](#)
- **explaining concepts at different levels of sophistication** for students who may benefit from additional context or from having complex concepts explained in different terms or with examples. AI tools that are integrated with search engines can be used to search for things and integrate those search results into its answer to the thing you actually want to know about (e.g., Search for X. Search for Y. Then use those searches to explain Z.). [Example](#).
- **summarizing texts at different levels of sophistication** has the generative AI provide students with a summary of an article or dataset at a level of sophistication appropriate to their learning level.
- **create examples (many, many examples)** in order to practice concepts or skills from the course. Generative AI tools can create examples from across different disciplines, at different (and increasing) levels of complexity. Students can then personalize the examples to their skill level and practice the key skills of the course.
- **making a study plan** and make suggestions for time and effort based with the specific context and goals of each student in mind. [Example](#).

Generating Academic Content

Invite students to use a generative AI tool to:

- **expand or condense text** (e.g., expand bullet points to actual text) or condense longer text into shorter text (e.g., condense text into bullet points). A related function is to use AI to summarize the key points from a text, including academic articles. [Example](#).
- **brainstorming / generating ideas** can be a useful starting point for students working to identify questions, topics, themes or arguments. A generative AI tool can also be asked to provide counter-arguments for a student-generated idea, that the student then needs to account for in their own work.
- **finding sources or references**. This is a capability where you have to be extra careful. As we know, generative AI tools can “hallucinate” sources that do not exist. Generative AI tools that are integrated in search engines generally perform better at this task. Regardless of the tool used, it’s good practice to verify that any sources identified actually exist. [Example](#).
- **identify and analyze data** from different datasets and **structure tables with information** based on inputted text or data samples with specific parameters offered [Example](#).
- **interact with spreadsheets** generative AI tools like ChatGPT can easily read the .csv format. You can extract a CSV file and give it to ChatGPT to work with based on certain specifications (e.g., give me an overview of what’s in this CSV file and provide some insights into the information provided), as well as output a CSV file. [Example](#).
- **coding with** natural language prompts, complete partially written code with suggestions, or translate code from one programming language to another. [Example](#).

With all of these uses it’s important to remind students that what the generative AI tool generates may have hallucinations or biases. Students should be reminded to review and evaluate the output from the generative AI tool to ensure its accuracy and evaluate its effectiveness.

You may be wondering – or your students may wonder – what generative AI tool to use for these tasks. [This review essay by Ethan Mollick](#) summarizes the capabilities of the major generative AI tools and makes suggestions on the best tool to use for a specific task. You can also visit [“There’s an AI for That”](#) to find new generative AI tools for specific educational tasks.

Student Perceptions of Generative AI in Teaching and Learning

Since the launch of ChatGPT in November 2022, there have been a few small-scale studies to investigate student use and perceptions of generative AI, many with the aim of describing different levels of experience and different perspectives among students. Unsurprisingly, students, like instructors, evince a [wide range of reactions and uses of generative AI](#). Our purpose here is to use the limited available understanding of student

use and perception to inform McMaster’s guidelines and resources for students, while also recognizing a need for further investigation and [partnership with students](#) to explore the nuances of student perceptions of generative AI.

Student Use of Generative AI

While there are now surveys of students from around the world (see [Hong Kong](#), [Australia](#), [Spain](#), and [United States](#)) there are only a [few Canadian post-secondary specific studies](#), including [this one](#) from February 2023 of 1039 students related to student use and perception of generative AI.

At the time of publication, across all reviewed surveys, the findings show that while most students are familiar with generative AI, and ChatGPT specifically, and many students are experimenting with generative AI, few students report using generative AI within academic courses – either with permission from instructors or without permission. Indeed, in the Canadian survey the analysts found that based on reported responses “there appears to be little reason to fear rampant ChatGPT-assisted cheating within our sector at this point” (Millian & Janzen, 2023).

Where students do report using generative AI, the surveyed results point to use related to exploratory investigation of the tools, explanations of course topics and as a “learning aid.” [One researcher](#) found students were using ChatGPT as a “search engine,” a particularly problematic use case given the propensity of ChatGPT to generate factual errors and to ‘hallucinate’ citations.

AI Literacy Among Students

This ‘search engine’ example of ineffective use of generative AI in a teaching and learning context points to one of the significant areas of common interest in work related to student use of generative AI: [proactive and intentional education for students](#) in the capabilities, limitations and ethical considerations of generative AI tools.

Both students and instructors have expressed a desire for [AI literacy training](#) for students, with students [emphasizing a need for this education](#) as a means of both using generative AI ethically and effectively in academic work and in preparation for future careers.

While introducing AI literacy within the academic setting (in addition to peer-to-peer and self-directed learning) will benefit students’ ability to use the tools effectively in academic and career settings, the academic integrity literature also suggests that if students are taught how to use generative AI tools effectively and in accordance with the expectations of the course, they are likely to use these tools appropriately.

Perceptions of Generative AI Among Students

That students – so far in the survey results – are not using generative AI for widespread academic dishonesty runs counter to the popular narrative in the news media (and some of our anecdotal conversations) of a wave of cheating among the ‘digitally fluent’ and ‘students-as-consumers’. Assessing the assumption that students will cheat requires careful consideration of a constellation of unjustified negative beliefs about students that many of us carry, and how these assumptions may defensively position students.

In the [reported surveys](#), students instead express a range of [thoughtful and measured reactions](#) to the expanded access of generative AI, including concerns about equality of access to the tools, risks of misinformation/inaccuracy, privacy concerns and fundamental questions about the implications of these tools for their future careers.

It is important to reiterate here the range of perceptions among students. Just as instructors – and those in our Task Force – represent a wide range of views on the possibilities and risks of generative AI, so too do students hold different perspectives. Included surveys point to differences among the age of students, disciplinary backgrounds, and domestic and international populations; the surveys also highlight differences in attitudes based on those who had tried generative AI tools and those who had not.

Recognizing this heterogeneity, the McMaster guidelines and resources for students begin from an assumption of different levels of experience with and exposure to the tools, different degrees of interest in using them and concerns about the impact of them on the teaching and learning environment.

Partnering with Students

Many universities are [communicating with students about generative AI in some manner](#), including reminding them of academic integrity policies, posting FAQs about LLMs, or creating citation guides. [Student disengagement](#), owing to trends that predate the pandemic but were exacerbated by the experience of remote learning, suggest that communications and AI-detection tools alone will be insufficient to respond to generative AI in education; instead, meaningful partnership and dialogue with students may help develop a response that privileges student learning and lead to more transformative and enduring solutions to the shortcomings of higher education, including student-led research, ongoing discussions and student-centred strategic planning, all of which go beyond the “token-student-representative-on-a-university-committee” model ([Abebe & Amarasinghe, 2023](#)). An early example of a student-centered response to policy development for generative AI comes from Boston University, where a Data, Science and Ethics undergraduate [class collectively wrote](#) a policy for generative AI that was then adopted as official departmental [policy](#).

Talking With Students About Generative AI

While we are working to share resources with students about using generative AI, and developing AI literacy materials, it's worth talking with your students about generative AI – especially if you plan to use it or have your students use it in the course. You might consider sharing:

1. **Definition and Use Cases:** Start with a brief overview of what generative AI is and how it works.
2. **Training Process:** Discuss how generative AI models are trained using large datasets and how the quality of output depends on the data it's trained on.
3. **Capabilities:** Touch on how these AI models can generate novel content, emulate human-like text, produce artworks, and how they're being used in diverse areas like entertainment, education, and research.
4. **Limitations:** Discuss the limitations such as inability to truly understand context or human emotions, factual errors or 'hallucinations,' the risk of generating inappropriate or biased content, and the challenge of generating long, coherent narratives.
5. **Ethical Considerations:** Discuss ethical issues such as potential misuse of AI-generated content for disinformation or deepfakes, copyright considerations, and privacy issues related to the use of personal data in training these models.
6. **Future of Generative AI:** Discuss the future potential of these tools, including the role they could play in society and the kind of regulations or policies that might be needed.

You then might want to have an open discussion with your students – full disclosure, ChatGPT helped generate these discussion questions! but we edited them.

1. How can generative AI be used to enhance teaching and learning in a post-secondary setting?
2. What are the potential drawbacks or risks of incorporating generative AI into the educational process?: This might stimulate a discussion on issues like fairness, privacy concerns, and the potential for AI to make errors.
3. Can AI-generated content replace human instructors for certain teaching activities? If so, which ones and why?: This question encourages students to consider the value and irreplaceability of human instruction.
4. How can we ensure that the use of AI in education doesn't exacerbate existing inequalities (such as access to technology or learning opportunities)?: This could spark a debate about equity and accessibility in education.
5. Should students be taught how to interact with and utilize AI as a part of their educational curriculum? Why or why not?: This can lead to a discussion on the importance of digital literacy in the 21st century.
6. How might generative AI impact academic integrity? Could it lead to increased plagiarism or other

forms of academic dishonesty?: This question can provoke thoughts on how technology might be misused in an academic setting.

7. How might the use of AI in formative assessment change the teacher-student relationship? Could it make the process more objective, or could it devalue human judgment and feedback?: This question can stir a discussion on the human aspects of teaching and learning.

6.

USING GENERATIVE AI AS AN INSTRUCTOR

While our chapter on student learning offers examples of when and how you might consider incorporating generative AI into your courses and assessments for student learning, this chapter focuses on ways you might use generative AI to support your teaching.

At McMaster the Provisional Guidelines on Generative AI in Teaching and Learning ask that if you do use generative AI in your teaching materials or assessment practices that you only share this with your students both in the course outline and in class. Sharing your use of generative AI with your students is intended to build trust and transparency, and to acknowledge that you are also using – and learning about – generative AI.

Those same Guidelines suggest that you can use generative AI with data collection turned off to provide students with formative feedback on assessments. Formative feedback is feedback that is not for grades, but rather gives students fast and specific advice on how to improve. Formative feedback from a generative AI tool might be given on an essay outline or draft, for instance, while you or the teaching assistant would be responsible for assessing and grading the final essay submission.

Finally, the Guidelines also ask that you check the accuracy of any AI created content. Recognizing that these tools “hallucinate” – or come up with factually incorrect responses – it is important that you check the accuracy of any content you might use in class, or any feedback offered to a student.

With that said, here are some broad categories where generative AI may be useful to you as an instructor:

Generating Test Questions and Assignments

By prompting a generative AI tool with the specific context of your course, as well as the subject you are aiming to assess and the kind of question or assignment you are interested in, the generative AI tool can offer many – many – examples of test questions at different levels of complexity, or different types of assignments. You can even ask for assignment ideas that meet the criteria of authentic assessment discussed in the chapter on assessment, or for assignment ideas that incorporate pedagogical approaches you value (e.g. problem based learning, community engaged learning or case based learning).

Generating Examples, Explanations and Counter Positions

Students benefit from practicing what they are learning with examples. Many, many examples. Generative AI

can be powerful in producing lots of examples for students to practice with, while also providing students with feedback on whether their submitted responses are correct, or how they might improve on a response. This personalized, immediate feedback is incredibly powerful for learning.

It can be challenging sometimes to describe a concept at many different levels of complexity. Some courses – especially those with no prerequisites – may have a range of experience and abilities in the class. Using generative AI tools you can quickly develop (and then check for accuracy) multiple explanations for a course concept. You could even have these explanations be written in unique and memorable ways – like, explain the carbon cycle in a limerick or describe the Canadian political parties as characters on the Simpsons.

Generative AI tools like ChatGPT can take on different personas by prompting – for instance, you could ask the tool to “pretend you are a heart surgeon” or “act like you are the Prime Minister”. In assigning this persona, the generative AI tool will produce text written as if from that position. This kind of role can be useful in inviting unique perspectives into a class discussion, or providing a provocative counter point.

Gathering Ideas for Class Activities and Assessments

Confronted with the challenge of generative AI you may be looking for new ways to teach a concept or skill, or new ways to assess a learning outcome. Generative AI can provide customized suggestions for interactive and engaging classroom activities (e.g. suggest six different interactive ways I could teach an auto-ethnographic research method to a third year, online class of 60 students in Sociology), as well as assessments that either incorporate generative AI or make generative AI less likely to be used.

For more examples of how you might use generative AI as an instructor, you can check out this newsletter, [One Useful Thing](#), by Ethan Mollick.