

Artificial Intelligence: Considerations for Graduate Research

Ontario Council on Graduate Studies

October 2024

Artificial Intelligence: Considerations for Graduate Research

Ontario Council on Graduate Studies

October 2024

This document will be revisited and updated in 2026.

Introduction

Generative artificial intelligence (AI) uses machine learning to generate new content in the form of text, audio, video, code and images. While AI is not new – researchers have been working with machine learning and deep learning for decades – the popularization of generative AI with the launch of ChatGPT in November 2022 has changed public awareness of its potential. Alongside growing awareness, there has also been an acceleration of developments in AI with new or improving large language models (LLMs) and associated tools rapidly being released.

The opportunities and challenges that generative AI tools present for graduate research are both complex and contextual. Differences in institutional approach, discipline, supervisor and student values, and research areas make broad recommendations challenging; likewise, different granting agencies and journals have or will have unique approaches to the responsible use of generative AI.

This document aims to provide an overview of the questions that graduate schools, graduate supervisors, and graduate students may have with respect to generative AI and graduate research, as well as some possible approaches to these questions. It focuses primarily on the graduate research enterprise, recognizing that in some instances these questions will apply to other aspects of graduate studies, including coursework. This document does not aim to be prescriptive or interpreted as policy; rather, it aims to stimulate local and contextual discussion and decisions by providing broad considerations.

What are some of the capabilities of generative AI and how could it be used in graduate research?

With thousands of tools, and additional capabilities emerging, the possibilities for the use of generative AI in graduate research are many. Some of the common capabilities of foundational models and tools like OpenAI's ChatGPT, Microsoft's Copilot, Anthropic's Claude, and Google's Gemini include summarizing, simulations, translation, outlining and brainstorming, data analysis and visualization, textual analysis, image generation and interpretation, captioning and transcription, video creation, copy editing, audio production and analysis, and more.

For graduate students these capabilities may be explored in different phases of graduate research from ideation to structuring research, to conducting research, to dissemination of findings, and to knowledge translation. Possible use cases might include using generative AI to identify research gaps in existing literature, framing research questions, developing literature

reviews, drafting qualitative or quantitative instruments, analyzing and interpreting data, drafting written or visual analysis, or shifting academic writing to plain language.

While these applications may benefit all graduate researchers, the capabilities of these tools may increase accessibility for graduate researchers with disabilities. For instance, the ability of a generative AI tool to provide a summary of a meeting transcript, to describe images and videos in text or spoken language, or to compose text, may increase accessibility. The potential for greater accessibility with generative AI tools continues to be explored.

Likewise, translation capabilities of these tools introduce opportunities for students for whom English is an additional language or for graduate research that benefits from translation as part of knowledge mobilization.

Given the sheer breadth of capability and the range of possible uses within graduate research, it is prudent for individual graduate programs and graduate supervisors to carefully consider what aspects of the program or individual field of study might benefit from the integration of generative AI, and what areas may not. While these conversations may currently take place between individual supervisors and graduate students, there will be a need for consistency across supervisors, programs and institutions. The following questions are intended to stimulate these individual conversations; OCGS remains attentive to communicating emerging norms.

How are academics approaching the use of generative AI in graduate studies so far?

Perceptions on the use and value of generative AI for graduate studies understandably vary, often significantly, among individuals, academic disciplinary cultures, and institutional settings.

In some disciplines, where machine learning has a longer history, there is broader understanding of the capabilities and limitations of AI, as well as recognition of the strengths of this technology for solving problems that were previously intractable.

In other disciplines, AI calls to the fore questions of ethical engagement and social and environmental responsibilities.

Just as there are disciplinary differences, there may be differences between how supervisors and graduate students approach AI. For supervisors there is a call to understand and evaluate the possible value of generative AI alongside their graduate students; for graduate students there is an invitation to discuss the same with their supervisors. For both supervisors and graduate students, the imperative to transparently discuss and share when generative AI will be or has been used is essential for both building trust and establishing norms of use.

Likewise, should a supervisor find opportunities to use generative AI tools in the practices of graduate supervision – providing formative feedback, summarizing relevant literature, etc. – a similar and shared expectation of transparency and citation of use is warranted.

In all areas of the university and across all roles, there is a need to weigh the potential value of responsibly used generative AI against its risks and challenges. Some of those risks and challenges include maintaining data security and privacy to protect intellectual property, hallucinations and bias in outputs, environmental impact of model development and tool use,

and questions of AI governance and policy at the national and international levels, including issues of copyright and AI safety. Understanding these risks and staying current on developments in AI technology is its own challenge.

For graduate research, in particular, the use of generative AI cannot compromise the unique contribution of the researcher as creating knowledge or its application. Indeed, the purpose of graduate research – the pursuit and creation of new knowledge through curiosity and critical engagement – cannot be matched by that of generative AI.

Given the capabilities of generative AI, should graduate degree level expectations or the values of graduate research change?

No. In keeping with the [OCGS Principles for Graduate Study at Ontario Universities](#), graduate students, both master's and doctoral, are assessed within their disciplines and against the degree level expectations of all graduate studies at Ontario universities and elsewhere, inclusive of depth and breadth of knowledge, conceptual and methodological awareness, communication skills, application of knowledge, professional capacity/autonomy, and awareness of limits of knowledge. Central to these expectations are the values of curiosity, integrity, transparency, and accountability.

While generative AI may change parts of the process of graduate research and may, over time, inflict these degree level expectations, these core expectations for all graduate students remain. Less concretely, but no less significantly, graduate research holds the reward of well-articulated knowledge advancement and innovation made possible by human ingenuity, curiosity and critical thinking. Novel contributions to research are required both for the betterment of our world, and – perhaps more cynically – for the further development and operation of the data sources on which generative AI tools are trained. If generative AI can enhance parts of the process of graduate research, the researcher remains accountable for, and ennobled by, each decision in the research process.

It is the expectation that if generative AI is used by graduate students, it will be in service of support or enhancement of the achievements of graduate research, but never as a substitute.

Given the capabilities of generative AI, should program learning outcomes, scope of graduate research, or graduate research activities change?

Yes. As generative AI tools and their capabilities are further explored and better understood within disciplinary settings, there stand to be shifted or expanded competencies and outcomes within graduate programs and for graduate researchers.

Just as each graduate program envisions the program learning outcomes, program structures to achieve those outcomes, and assessments to ensure the successful completion of the outcomes, each program will find occasion – perhaps through the institutional quality assurance process – to reexamine the goals of the program in light of the capabilities of generative AI. For instance, common outcomes related to completion of a literature review might be expanded to include critical appraisal of an AI-produced literature review.

What are the implications of generative AI for graduate research and academic integrity?

Institutions have different academic integrity policies and processes, just as different institutions and disciplines are developing norms and expectations on referencing, citation and acknowledgement of the use of generative AI tools in academic work. Nevertheless, there are core principles of academic integrity that apply across all our academic endeavors and institutions: integrity, information security, transparency, accountability, ownership, and responsibility. Graduate students as individuals and as authors of the thesis have ultimate responsibility for understanding and adhering to the academic integrity policies and expectations of their institution. Graduate programs have the responsibility to develop and implement assessments of graduate students that confer confidence in the capabilities and competencies of their graduates (see evaluation below).

Wherever possible, supervisors and supervisory committees should provide clear guidance and expectations on if and how generative AI can be used by the graduate student in the research process and in the dissemination of findings. These expectations should also include how the graduate student should transparently declare and acknowledge when and how generative AI tools were used. It is worth emphasizing that the norms of referencing are evolving, and that different disciplines and institutions are approaching referencing practices with some variation.

What are the implications for evaluating the attainment of program learning outcomes and degree level expectations?

Just as conversations about evaluation are evolving at the undergraduate level and in course-based master's programs, there will be need for reflection and revision to evaluation practices across graduate research programs.

When a graduate researcher's oral defense is most successful it demonstrates a value at the heart of the research enterprise: peer review. That is to say, the graduate researcher's oral defense is an occasion both to safeguard and validate the student's competencies and capabilities, but is also an opportunity to gather academic peers for thoughtful engagement with the research, to probe and test the research, enter into the discussion with openness to new ideas, to contemplate alternative perspectives and frameworks, and to push the boundaries of what is already known. While the candidacy/comprehensive/qualifying exam and oral defense is a means of ensuring academic integrity, its primary purpose ought always to be the demonstration of this rich tradition of peer engagement and scholarly inquiry.

With that, graduate programs and units are encouraged to engage in collective conversations about acceptable uses of generative AI across the discipline or program, as well as how to reference such use in policies and processes. Each institution will have different means of enabling and supporting these conversations; the quality assurance process may be one possible means of holding such discussions.

What comes next?

The Ontario Council on Graduate Studies (OCGS) will continue to discuss the impact of generative AI on graduate research. Some known areas of further work include the

development of resources for graduate students on generative AI, as well as resources for oral examination committees. Suggestions or comments on this document, or additional areas of support are welcome and can be directed to Katarina Todić, Senior Policy Analyst, COU (ktodic@cou.ca).

Conversation Starters

Below are some questions that supervisors and graduate students may discuss in deciding when and how generative AI could be used in graduate studies. These questions are framed as institutionally and disciplinarily agnostic and may be adapted depending on disciplinary and institutional contexts. Likewise, these questions are not intended to be definitive or to limit other areas of consideration that may be relevant to an institution or the student's research. Finally, these conversations about AI use could occur more than once as graduate research takes place and as advances in the technology unfolds.

Supervisor and Graduate Student Conversation Guide

- 1. What do you [the graduate student/the graduate supervisor] already know about generative AI and what might you need to learn before incorporating these tools into your graduate research? What is your individual approach to generative AI? What do you believe about its value or risks?**

Possible discussion prompts:

- How would you describe your 'philosophy of AI use'? When, how, and why do you think AI should be used in research?
- How might generative AI intersect with, influence, or impact your professional goals?

- 2. What scholarly activities within my graduate research may benefit from the use of generative AI? How could it enhance my ability to achieve the [degree level expectations](#) and program learning outcomes?**

Possible discussion prompts:

- What value might generative AI bring?
- What impact might using generative AI for [this task] have on my core learning experience as a graduate student?
- What phases of the research process and research activities would most benefit from the inclusion of generative AI?

Examples: translation, copy-editing, brainstorming, concept explanation, drafting, coding, data analysis, data visualization, drafting, simulations, literature reviews

- 3. What scholarly activities within my graduate research should not involve the use of generative AI? In which cases would the use of generative AI lessen or preclude my ability to achieve the degree level expectations and program learning outcomes?**

Possible discussion prompts:

- What might be some of the risks of using generative AI to complete [this task] or [this part]?
- What could be some of the negative impacts on my studies, colleagues or my disciplinary community if generative AI was used for [this part]?
- What impact might using generative AI for [this task] have on my core learning experience as a graduate student?

Examples: translation, drafting, data analysis, data visualization, interpretation and analysis, synthesis, literature reviews

4. What benefit or risk does the use of generative AI pose for me as a graduate researcher?

Possible discussion prompts:

- What generative AI tool(s) will I use and what risks or benefits does that tool offer for data security and privacy?

Examples: accessibility features, data sovereignty, implicit bias, data protection, privacy, [data contamination](#)

5. How should I document and disclose when I have used generative AI in my work? What level of use (e.g. brainstorming, drafting, copy editing, coding) warrants disclosure of use? How do I ensure everyone involved in the work I am doing understands how we will use (or not use) generative AI? What are the institutional expectations to consider?

Possible supervisor strategies:

- Consult with institutional services and supports to find existing guidance and/or policies that may apply.
- Citation and disclosure practices vary by context. Check with colleagues, journals and funding agencies in your disciplinary area to consider what emerging norms for citation or disclosure may be.
- Sample acknowledgement could read: “[Name of generative AI tool] was used in the creation/drafting/editing of this [scholarly output]. I have evaluated this document for accuracy.”

Possible discussion prompts:

- What research ethics implications and obligations do we have to consider?
- What might be some reasons our [key consulted groups] might need or want to be aware that generative AI was used in this [type of work]?

Artificial Intelligence: Considerations for Graduate Research

- How do we ensure that everyone involved in a project or process that uses generative AI is aware and agrees to the use?
- What possible risks to our credibility or expertise are present if we do not disclose use of generative AI in this [type of work]?
- What professional obligations do we have to be transparent with our use of generative AI in our area?

Supervisors who may want to practice this conversation ahead of time and/or receive feedback on their responses may use the following prompt in a generative AI tool of their choice to simulate the conversation and receive feedback:

- “I am a graduate supervisor in [insert your Department] at [insert your university.] My university wants all graduate supervisors to have conversations with their graduate students about the use of generative AI in the research process. They have provided the attached guidelines [attach these guidelines], including conversation starters and discussion prompts. I want to practice having this discussion with my student before I meet them. You will play the role of my graduate student, and I will play the role of supervisor. You will respond to my questions with plausible answers and make up any details related to your research that will enable you to provide nuanced and sophisticated responses. After each of my responses give me feedback on what I could say to improve my response or more information I could provide. After providing me feedback, continue with the simulation.”



For more information, please contact:

Council of Ontario Universities
180 Dundas Street West, Suite 1800
Toronto, Ontario M5G 1Z8
416-979-2165
cou@cou.ca

Connect with us:

[COU website](#)
[Like us on Facebook](#)
[Follow us on Twitter](#)
[Follow us on LinkedIn](#)