McGill | Graduate and Postdoctoral Studies

Using artificial intelligence (AI) in research writing – guidelines for graduate students

Purpose

The purpose of this document is to provide **clear guidance** for students and instructors about how and when generative AI tools (for example, ChatGPT) can be used in graduate studies and research, as per UNESCO recommendations.

Introduction

Generative artificial intelligence (AI) tools have become commonplace in many domains. These tools allow for users to supply prompts to a computer algorithm that generates automated responses based on a repository of accumulated information and data. The AI-generated responses come from an underlying probabilistic model that has been trained by data coming from public and private online data sources, as well as human input as to the usefulness of past responses. One of the early generative AI tools is OpenAI's ChatGPT, although newer tools[1] that provide similar or expanded functionalities also fall under these same guidelines.

The class of Large Language Models (LLM's) is the most widely used class of models for generative AI. LLM's provide a structured framework to quickly respond to complex user prompts with responses that are often not easily distinguished from the responses one would receive from human experts. This differentiates tools that use LLM's from other tools which provide simple, targeted suggestions to input, such as auto-correct in word processing applications.

With respect to research, generative AI tools have seen wide application in programming (both for debugging previously written code as well as generating new code from user prompts) and are now gaining traction in the scientific writing domain. Beyond more traditional language applications, they have also been used for generating other material, such as music or images.

Understanding the ethics and limitations of using generative AI in research writing

To understand where and when graduate students should utilize AI tools in their work, they must first consider why they are doing the work in the first place. At the graduate level, students are expected to not just learn things and to do things, but to learn *how* to do things and to become *expert* in the things they are doing. Equally important is the ability to critically analyze information and make decisions based on those analyses. As part of their training, graduate students need to learn and adopt good research practices to eventually produce original research.

If one uses generative AI to generate a complete abstract based off vague bullet points, or if one asks a question and uses the response as a paragraph in one's work, one would not have engaged in the research **process**; that is, one would not be able to repeat the writing process on one's own without

[1] There exists a wide range of generative AI tools other than ChatGPT, such as Research Rabbit, Consensus, Perplexity, Elicit, Quillbot, and more.

generative AI. This would be analogous to a person trying to learn how to drive by standing on the sidewalk watching another person drive a car. Not only is the person on the sidewalk still incapable of driving the car themselves, but they would also not even know if the person driving the car was doing it correctly. The person on the sidewalk is not engaged in the **process** of driving the car, only in viewing the result, a result which is difficult to understand and learn from due to their own lack of experience. Watching someone else exercise will similarly not increase one's own level of fitness. Learning to do research requires engaging and even struggling with difficult and sometimes frustrating tasks, so outsourcing these tasks to a generative AI tool can inhibit one's learning.

The increasing prevalence of generative AI tools in university settings has also led to concerns related to academic integrity. This is especially true in the context of assessment and learning outcomes, and in the creation of original research. While generative AI can be used as an intelligent assistant to save time (e.g. when improving the quality of one's writing or gathering information about a new topic), it can also lead to unintentional instances of plagiarism and cheating.

Assume that one wants to use a generative AI tool, e.g. ChatGPT, to quickly create an abstract or even to generate complete paragraphs of a research paper. AI tools produce output based on models built from existing written material and user feedback, the vast majority of which will not be specific to one's writing goal. Therefore, although the AI tools can do this sort of work, the text that is produced can frequently suffer from one or more of several potential problems:

- **Plagiarism**: In response to a prompt, ChatGPT may inadvertently reproduce exact text from the other sources without attribution. (Sun & Hoelscher, 2023) Even if it does not produce exact text from other sources, it is an open ethical question as to whether definitions of plagiarism should include text produced by AI tools that *sounds* like it was written by another person because the models are based on that person's body of work. (Lund et al., 2023)

- **Falsification**: When ChatGPT does not exactly reproduce existing material, it can sometimes produce new statements that are false. For example, it may not correctly state complex facts, even though a human would immediately realize they were incorrect. In some cases, ChatGPT can even generate falsified research references, that is, plausible citations with reasonable title, author, and journal combinations for papers that do not exist. (Buriak et al., 2023; Sun & Hoelscher, 2023; Zheng & Zhan, 2023)

- **Wrong target audience**: If generative AI is not used properly (and even sometimes when it is), one can generate text that is correct, but that is not appropriate for the audience to which one is writing or the genre in which one is working. In other words, the language may be too colloquial for an expert reader or too formal and technical for a lay reader. The generated text may contain words or phrases that are not used by members of the community to whom one is writing. Generative AI tools can very much struggle with these *context* issues that are easily detected by human readers who work in the domain of your research.

- **Nonsensical results**: ChatGPT, especially when asked to produce large sections of text, can sometimes give contradictory statements or produce ridiculous statements that are perfectly grammatical, but do not make logical sense. (Barrot, 2023; Sun & Hoelscher, 2023; Zheng & Zhan, 2023)

- **Loss of voice**: In the end, if one does not write the text, one's personal message and voice will be lost. The less the text sounds like the writer, the less ownership the writer has over the ideas and sentiments communicated to the reader. Most research on AI-generated text refers to a lack of originality in much of the produced work, particularly when ChatGPT is given broad instructions or a large body of text to write at once. (Buriak et al., 2023)

- **Ethical violations**: The generative AI tools generate text based on a wide variety of sources which are unknown, and it cannot make users aware of the specific source. They may use information that is supposed to be private (especially if that information was irresponsibly or inadvertently entered into a prompt by another user). Without appropriate sourcing, one may be complicit in and even exacerbating violations of privacy or intellectual property rights.

Users of generative AI should be aware of the above limitations, as well as additional restrictions from real-world entities (e.g. some journals have outlawed the submission of any text that is predominantly AI-generated; McGill requires that student's submit work that is their own work, so content submitted that is completely generated by generative AI clearly violates the McGill Code of Student Conduct).However, even work that is not completely generated, but is a mix or derived from generative AI could also result in potential violations. Furthermore, users may also want to consider the harmful impacts of generative AI on the environment.

We encourage supervisors and instructors to set clear expectations of how generative AI will be used in research and in coursework by students.

- Being accountable and transparent about your work
- Investigating when using generative AI is appropriate
- Conducting research with academic integrity
- Reflecting on when/why generative AI is being used, and refraining from using it to replace learning

Students' responsibilities related to the use of generative AI tools for research

Evaluating the potential benefits and risks of certain generative AI usage in graduate studies

Graduate students **can** use AI to help them engage more efficiently in the **process** of learning how to do research, when used wisely. Below, the different uses of generative AI have been divided into three categories: green, yellow, and red. The examples listed in the green category constitute acceptable uses of generative AI in graduate studies, which focus on using the tool as an aid in doing the work rather than to eliminate important steps of the process. The yellow category calls for prudence and requires careful consideration of implications, while the red category covers use of AI that is not permissible. Simply put, the green examples *support* the work the student is already doing (e.g. by improving the quality of writing), whereas the yellow and red examples end up *producing* work the student is expected to do themselves. Generative AI should not be used to replace that work and in those situations, it will prevent learning, rather than support the learning process.

Green: using generative AI as a tool

Definition: Using generative AI for proofreading and editing existing text (one's own work) is generally accepted. This can be particularly helpful for students whose first language is not English. In these instances, the tool is used to improve text, not produce it. Generative AI tools can also be helpful as search engines when looking for existing literature on a given topic, although remember to verify any references, as some may not be accurate or even exist.

Examples of the use of generative AI in this category include:

- Proofreading: Spell check, grammar check
- Editing: Lightly rephrasing objectives in grant writing or tightening language in a draft paragraph, although one should be cautious of not losing one's voice
- Translation: translation of existing text to another language (e.g. translating one's English abstract into French), although this is dependent on the field of study and the purpose of the translation
- Programming: Debugging existing code
- Search engines: Using generative AI to sift through and identify relevant research articles for a project
- References: Formatting the reference list

Yellow: using generative AI to generate content for one's own use

Definition: Using generative AI to generate content, even if restricted to one's own use, means producing something new. The student is responsible for the use of that content even if they don't claim it as their own and use it only for their own reference. It is worth noting that content produced by generative AI may generate falsified information or produce false references. When generating images and sounds, parts of or entire existing images or sound files may be reproduced exactly in ways that would be impossible to appropriately reference or provide sources for. Therefore, students can use these for their own personal use, but should be very wary of sharing them with others or including them in submitted or published work.

Examples of the use of generative AI in this category include:

- Learning: Using generative AI to break down complex topics or processes, for one's learning purposes
- Literature review: Producing summaries of published research on a given topic for one's own consumption
- Programming: Writing programs or generating code for one's own use that will not be distributed widely
- Research writing (grants, papers, publications): Generating suggestions for potential reviewers (suggested names may not be the most relevant for a given topic)
- Research design: Suggesting additional steps/tasks for experimental design (e.g. using generative AI to identify steps for one to consider, adapt and/or modify)

Red: using generative AI to generate content and presenting it as one's own work

Definition: Using generative AI to generate content and presenting it as one's own work constitutes plagiarism, in violation of the Code of Student Conduct. Work submitted by students for evaluation is expected to be their original work unless clearly indicated otherwise. Content generated by AI does not constitute original work.

Examples of the use of generative AI in this category include:

- Literature review: Submitting summaries of published research on a given topic as one's own work
- Programming: Sharing entire programs or large amounts of code produced by generative AI without proper attribution
- Research design/writing: Generating objectives, arguments, perspectives, research ideas, figures/images, etc.
- Research design/writing: Conducting analyses
- Research writing: Writing entire sections or the entire proposal/paper
- Research writing: Generating research quotations

Conclusion

Generative AI frameworks can be powerful tools which can increase the quality of one's research writing. Strong arguments have been made that prompt engineering will become an important research skill. By using specific prompts and instructions, generative AI tools can detect errors and provide suggestions for even subtle improvements. However, using generative AI tools to create new text can be quite problematic, and any output used from these tools should be thoroughly vetted to identify possible issues such as plagiarism, falsification or other negative side effects of using AI-generated text in one's work. In the end, at McGill, we want our students to learn to drive the car and not to simply stand watching on the sidewalk.

References

Barrot, J. S. (2023). Using ChatGPT for second language writing: Pitfalls and potentials. *Assessing Writing*, *57*, 100745. https://doi.org/10.1016/j.asw.2023.100745

Buriak, J. M., Akinwande, D., Artzi, N., Brinker, C. J., Burrows, C., Chan, W. C. W., Chen, C., Chen, X., Chhowalla, M., Chi, L., Chueh, W., Crudden, C. M., Di Carlo, D., Glotzer, S. C., Hersam, M. C., Ho, D., Hu, T. Y., Huang, J., Javey, A., … Ye, J. (2023). Best Practices for Using AI When Writing Scientific Manuscripts. *ACS Nano*, *17*(5), 4091–4093. https://doi.org/10.1021/acsnano.3c01544

Huang, J., & Tan, M. (2023). The role of ChatGPT in scientific communication: Writing better scientific review articles. *American Journal of Cancer Research*, *13*(4), 1148–1154.

Lund, B. D., Wang, T., Mannuru, N. R., Nie, B., Shimray, S., & Wang, Z. (2023). ChatGPT and a new academic reality: Artificial Intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology*, *74*(5), 570–581. https://doi.org/10.1002/asi.24750

Sun, G. H., & Hoelscher, S. H. (2023). The ChatGPT Storm and What Faculty Can Do. *Nurse Educator*, *48*(3), 119. https://doi.org/10.1097/NNE.0000000000001390

Zheng, H., & Zhan, H. (2023). ChatGPT in Scientific Writing: A Cautionary Tale. *The American Journal of Medicine*, *0*(0). https://doi.org/10.1016/j.amjmed.2023.02.011