LAURIER
WILFRID LAURIER UNIVERSITY

Use of generative AI in research

A sector scan and review of university guidance

Prepared by:
WLU Generative AI in Research Working Group

Version 5.1: April 8, 2024

Disclaimer

Rather than starting from scratch, we looked to the advice already provided on this topic by other universities and publishers. We then tailored this text to our needs by editing, paraphrasing, combining, and adding to it. We have credited the source of all material in this document, but for the sake of readability, we have not quoted it verbatim. Nevertheless, the originality of much of the content belongs to the original sources.

The field of consumer-available generative AI is rapidly developing! We have tried to provide the most up-to-date information, but in this field it will always behoove the reader to check the most current advice and information available before taking decisions or actions on the use of generative AI in their research.

Working Group

Azam Asilian Bidgoli, Taylor Coleman, Una Glisic, Mark Humphries, Jonathan Newman, Charity Parr-Vasquez, Ellen Russell, and Brent Wolfe.

With help from: Jennifer Ferfolja, Scott Gillies, Lianne Leddy, Audra Mitchell, and Nela Petkovic

Contact

Jonathan Newman
Vice President: Research
Office of Research Services

Contact: research@wlu.ca

Changelog

v1.0	2023-11-19	Completed first draft
v2.0	2024-02-16	Committee edits, additions, and deletions Incorporated; index and recommendations added
v2.5	2024-02-26	Committee edits, additions, and deletions. Reorganization of Section 5
v3.0	2024-03-04	Final committee edits before consultation process.
v4.0	2024-03-25	Incorporated comments from the Laurier community.
v4.1	2024-03-26	Added subsection 2.7 Generative AI detectors are not reliable
v4.5	2024-03-26	Added revisions related to EDI and benefits of gen AI
v5.0	2024-03-27	Minor edits. Final version for circulation
v5.1	2024-04-08	Minor edits

Table of Contents

1 In a nutshell: generative AI in research

Generative AI tools produce outputs whose origins are often unknown to users. Therefore, these tools may generate content that cannot be verified through primary sources. The output is based on pre-existing data, which may contain biases and other restrictions. It is important to identify and take into account any such limitations. It is also not guaranteed that generative AI tools comply with laws and regulations designed to protect confidential information such as Ontario's Freedom of Information Protection of Privacy Act (FIPPA) and the Personal Health Information Privacy Act (PHIPA).

Uploading data to a public AI tool is equivalent to making it public. Hence, it is vital to take appropriate steps to ensure that the disclosure of the information complies with all relevant rules and laws. Generative AI may also present other privacy risks, such as data breaches, exposure of intellectual property, and security issues.

Information: Generative AI refers to deep-learning models that can generate high-quality text, images, and other content based on the data they were trained on [1].

See *Top 20 Generative AI Tools & Applications in 2024* for a description of commonly used generative AI tools.

There are also intellectual property issues related to the ownership of AI-generated output. It raises questions about whether it should be treated as a published resource. The norms and requirements for citing AI-generated output and disclosing the use of AI technologies are complex, changing quickly, and often unclear. These norms and requirements may vary depending on the source. For example, publishers, journals, professional organizations, and funding organizations all have (or are developing) policies on the use of AI that must be followed [2].

This document has been created to provide information and guidelines on the use of generative AI in the research enterprise at Wilfrid Laurier University. The aim of this document is to incorporate materials and advice from various sources while also promoting responsible and transparent use of generative AI. It must be noted that this field is rapidly changing and perspectives may vary between individuals and disciplines. ***Therefore, this document serves as a set of best practices rather than a policy***.

2 What are some benefits of using generative AI in research?

The material in this section comes from an excellent article in Forbes Magazine by Beata Jones titled "How Generative AI Tools Help Transform Academic Research" [3].

AI-powered tools have made literature reviews and research more efficient and effective, allowing *experienced* scholars to focus on analysis and interpretation rather than searching for background information. Although there are many tools available to researchers, they vary in quality. Nevertheless, they give us a glimpse of what the future holds.

Services like Elicit, Consensus, and Inciteful help researchers manage a personalized collection of sources, allowing them to focus on analysis rather than continuously searching for background infor-

mation. Tools such as TLDRthis, AskYourPDF, and Scholarcy create summaries, extract key points, and more for lengthy papers, making them easier to understand. Powerdrill helps identify gaps in current literature, while Litmaps visually maps connections between various papers to help identify derivative works. Platforms like SCISPACE, Jenni, and MirrorThink can create an initial review of the literature. Julius and ChatGPT-4's Advanced Data Analysis can assist with data analysis. HeyScience provides feedback to help scholars align their work with academic standards. AI-powered editors, such as Grammarly and ChatGPT refine manuscripts by providing instant corrections. They can also format references. Generative AI can improve the clarity, style, and coherence of scholarly writing, benefiting non-native English-speaking scholars and individuals with disabilities [4].

At its (current) best, generative AI performs many laborious tasks and frees researchers to focuses on the more nuanced and specialised judgements that cannot be successfully managed by AI alone. However, generative AI also poses specific challenges related to transparency, understanding system results, cognitive barriers, incorrect information dissemination, manipulation, privacy concerns, and ownership.

AI algorithms and systems may help support and accommodate disabilities. They augment assistive technologies, robotics, and create personalized learning and healthcare solutions. Generative AI and language-based models further expand this impact and the R&D behind it. These systems may fuel existing assistive systems, health, work, learning, and accommodation solutions. They require communication and interaction with the patient or student, social and emotional intelligence, and feedback. Such solutions are frequently used in areas involving cognitive impairments, mental health, autism, dyslexia, attention deficit disorder, and impaired emotion recognition, which rely heavily on language models and interaction [5].

3 Principles and considerations for the use of generative AI in research

3.1 FASTER principles

We recommend that researchers adopt the FASTER principles [6], as recommended by the Government of Canada for the use of generative AI in government work. These principles also seem well suited for research applications:

Recommendation: Researchers should follow the FASTER principles.

- **Fair**: ensure that content from these tools does not include or amplify biases and that it complies with human rights, accessibility, and procedural and substantive fairness obligations.
- **Accountable**: take responsibility for the content generated by these tools. This includes making sure it is factual, legal, ethical, and compliant with the terms of use.
- **Secure**: ensure that the infrastructure and tools are appropriate the information and that privacy and personal information are protected.
- **Transparent**: identify content that has been produced using generative AI; notify users that they are interacting with an AI tool; document decisions and be able to provide explanations if tools are used to support decision-making.

- **Educated**: learn about the strengths, limitations and responsible use of the tools; learn how to create effective prompts and to identify potential weaknesses in the outputs.
- **Relevant**: make sure the use of generative AI tools supports user and organizational needs and contributes to improved outcomes for Canadians; identify appropriate tools for the task; AI tools aren't the best choice in every situation.

3.2 Privacy and data confidentiality

Inputting private or confidential data into a public AI system is the same as making it public. Submitting information to public AI tools, such as by entering questions or requests into tools like ChatGPT, is a way of releasing those data to a third party. *The same regulations that apply to other forms of public disclosure of private or confidential information also apply to interactions with public generative AI tools*. Sharing this type of information with public AI tools puts both individuals and the institution at risk of privacy and security violations. Similarly, uploading research data, grant proposals, or analytical results into a public AI tool is the same as making them public [2]. The items you upload are used to further train the model, and may show up as responses for another user. It is not recommended to share personal information with AI chatbots due to potential risks. OpenAI stores conversations on its servers and may share these data with third-party groups. This was demonstrated in March 2023 when a security breach caused some ChatGPT users to see conversation headings that did not belong to them. This incident was taken seriously, as ChatGPT had 100 million active monthly users at the time. In response, OpenAI made several changes, such as introducing an age restriction and making its Privacy Policy more visible. Additionally, it provided an opt-out Google form for users to exclude their data from its training or completely delete ChatGPT history. Even if you think you would not share your personal details, it is still possible to make a mistake, as demonstrated by a Samsung employee who shared company information with ChatGPT [8, 9].

Information: Public AI refers to openly accessible AI algorithms trained on diverse datasets, sourced from users or clients. ChatGPT is an example of public AI, trained on openly available data from various online sources. Providers of public AI refine their models by incorporating customer data, which may not be entirely confidential. [7].

Warning: Uploading research data, grant proposals, or analytical results into a public AI tool is the same as making them public

Recommendation: Supervisors should initiate conversations with their teams to ensure that all members understand both the opportunities and responsibilities related to the use of generative AI.

3.3 Transparency, attribution, and disciplinary norms

When using generative AI, researchers must stay up-to-date with the evolving standards and comply with the rules and norms in their research field. They are responsible for ensuring the accuracy of information, acknowledging authorship, using authentic sources, and transparently declaring the incorporation of AI in their research activities. Supervisors should initiate conversations with their teams to ensure everyone understands the opportunities and responsibilities related

Information: More detailed advice on citation conventions is provided by the Library on their Citing Use of AI (Artificial Intelligence) or Chatbots webpage.

to using generative AI. Researchers must also openly and transparently disclose and document their use of generative AI. They should acknowledge the use of generative AI software tools in their work and disclose the amount of new content generated. If entire sections of the work were created by these tools, an appendix or supplementary material should explain the use, including the specific tools and versions used, the text of the prompts provided as input, and any post-generation editing.[2]

3.4 What about 'embedded' AI tools?

Numerous software applications are incorporating generative AI tools or have plans to do so in the future. For instance, Microsoft's Word and Excel already have access to Microsoft Copilot, Microsoft's AI tool [10]. Note that Microsoft has also released a new version of its Bing search engine, confusingly also called Copilot, which functions similarly to ChatGPT and is a different product from its AI assistant by the same name. For more information from ICT, see Introducing Microsoft Copilot and Microsoft Copilot is now available!. Follow these pages for further updates.

Recommendation: We recommend erring on the side of caution and declaring any use of generative AI until these matters are more settled.

Some advice suggests that you do not need to disclose that you are using AI software tools to modify and enhance the quality of your text [11]. The logic being that it is similar to how you use a typing assistant, such as Grammarly, to improve spelling, grammar, punctuation, clarity, engagement, or a basic word processor to correct grammar or spelling errors [2]. However, we recommend erring on the side of caution and declaring any use of generative AI until these matters are more settled.

3.5 External guidance and policy

Researchers who wish to use generative AI tools have a responsibility to understand and comply with the policies and guidelines provided by journals, publishers, professional associations/societies, and funders/sponsors regarding the use of these tools for creating new content. While this document will be periodically updated, for active users of these tools, it is crucial to stay informed about the latest developments in generative AI. The technical capabilities of generative AI tools, as well as the regulations and standards related to their use, are constantly changing. Responsible research requires a current understanding of changes in AI technology and best practices within particular areas of research and scholarship. All those involved in research should strive to stay up-to-date with emerging AI tools, research studies, and ethical guidelines, and take advantage of professional development opportunities to improve their AI integration skills. Please see [2] and section 5 for more information.

Those conducting research with Indigenous or community-based research partners should take extra precautions by engaging in frank, open conversations about the benefits and costs of using generative AI in research, and to seek their partners' guidance on its use when working with them.

3.6 Generative AI tools have bias baked into them

Generative AI models reflect the biases present in the real world, as they are designed according to, and trained with, data that reflect existing inequalities and forms of discrimination, including racism, sexism, homophobia, ableism, transphobia, etc. This can result in discriminatory responses against various groups. Organizations developing AI models recognize the problem and are making efforts to mitigate it. However, the release of biased AI models to the public raises ethical concerns. The challenge with Large Language Models (LLMs) is that they may include content with biases, potentially perpetuating them further. Therefore, more comprehensive and ethical approaches to the development and deployment of AI models are needed [9, 12].

3.7 Generative AI detectors are not reliable

The primary issue with existing AI detectors is their insufficient level of accuracy. They often exhibit high rates of false positives and false negatives. While these systems can provide an approximate probability that the content is generated by AI, they are unable to directly verify the actual source of the text [13].

Present AI detectors give the impression of accurately recognizing artificial text, despite substantial evidence indicating otherwise. There are no definitive textual cues that consistently differentiate between human writing and sophisticated AI-generated content. Essentially, distinguishing between human writing and advanced AI is challenging as alterations in language, syntax, and coherence can deceive detectors. Their judgments are based on surface-level indicators and statistical data, rather than directly attributing the content to an AI model [13].

4 Use of generative AI in research with human participants

4.1 Overview

The emergence of AI has opened up a new realm of possibilities for researchers to integrate it into their studies. However, there is currently no consensus on the best practices for using AI in research involving human participants. Therefore, this guidance is subject to change as technology advances and the most appropriate methods for using it are established. This guidance covers the use of AI, machine learning, deep learning, and other related techniques in research activities, as well as other activities that may influence or regulate the use of AI tools in research [14]. Note that the definitive guidance will come from the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (TCPS 2), when it is updated.

4.2 AI use in research with human participants

When conducting research, a researcher may involve an AI tool in human participant research when: the AI tool collects data from humans through interaction or intervention; it is used to obtain, analyze, or otherwise access data about human research participants; or it acts as an extension or representative of the investigator(s) by answering questions for human research participants. As an example, tools providing transcription services are used in some human participant-based research, and some of these tools use AI. The following solutions can be utilized for transcription services at Laurier. These solutions have been reviewed through the Privacy and Security Impact Assessment (PSIA) process and approved for specific use cases and data types. Additional information about the data classes is available in Policy 9.5 External Information Technology and Cloud Services. As per Policy 9.5, approval through a Privacy and Security Impact Assessment (PSIA) is required before using any software tool with Type 2 or Type 3 data.

- Zoom: Approved for specific use cases with Type 2: Internal Information.
- Zoom for Healthcare: Approved for specific use cases with Type 3: Restricted Information.
- Microsoft Teams: Approved for specific use cases with Type 3: Restricted Information.
- Microsoft Teams: Approved for specific use cases with Type 3: Restricted Information.
- Microsoft Word and OneNote transcription service: Approved for Type 3: Restricted Information.

See Policy 9.5 External Information Technology and Cloud Services for definitions of data types.

The data to which other AI tools have access should be carefully safeguarded. Unless explicitly stated otherwise by the terms of use of the AI tool provider, there is no assurance that the information provided to the AI tools will remain confidential. Therefore, investigators should be cautious and avoid sharing any sensitive, identifying or private information when using these tools, particularly data that are legally protected. More broadly, the Research Ethics Board (REB) recommends that the collection of identifiable participant data be kept to the minimum necessary for the completion of the research activity, whether it is used with an AI tool or not [14].

4.3 REB review

Whenever an AI tool is used to interact with or acquire data generated by or from human research participants, it is considered to be engaged in human participant research and therefore requires review by the REB. The REB will assess the use of the AI tool in accordance with applicable ethical and regulatory standards, in addition to university policies. The REB will follow its standard procedures for the ethical review of projects, including delegated or full board review, and, when necessary, review from ad hoc or expert reviewers, in cases where the REB does not have the required expertise. For all research, the REB will review the use of the AI tool in the same context as any other research mechanism. This consists of minimizing risks to participants (including risks relating to privacy and confidentiality), free and informed consent, fairness and equity in research participation, and adequate protection of those living in vulnerable circumstances. In all cases, the REB will review the use of the AI tool in the same context as any other research tool, including the requirement that researchers use only University-approved software solutions to collect, analyze or store participant data [14].

4.4 Informed consent

The REB recommends, and in many cases requires, that investigators inform participants of the use and context of AI in all consent documents. This explanation should be in lay terms and should include the data to which the AI tool will have access, any limitations that will be placed on the AI tool and participants, and a description of what the AI tool will do with any data it receives. Investigators must also inform participants if their data cannot be removed from the AI tool. The REB does not permit AI tools to obtain informed consent from participants; a human investigator must be present. The REB will provide further guidance in due course.

Electronic consent processes, such as the use of Qualtrics, are allowed so long as the consent process is aligned with guidance outlined in the Tri-Council Policy Statement: Ethical Conduct for Research Involving Humans (TCPS2). AI tools should not have access to informed consent documents if an individual declines to consent. Researchers should have mechanisms in place to ensure consent is obtained from participants and that consent documents are separated from any research data being provided to the AI tool [14].

Recommendation: Inform participants, in lay terms, which data the AI tool will have access, any limitations that will be placed on the AI tool and participants, and a description of what the AI tool will do with any data it receives.

4.5 Identifiable data

Researchers should be increasingly mindful of the potential for AI tools to generate indirect identifiers in a dataset. The REB suggests that, where possible, investigators restrict the AI tool's access to demographic information and other data points that could be used to identify a person (e.g., while job title, employer, age and gender are not directly identifying data, together they may identify a participant). The REB recommends, and in certain cases requires, that investigators place restrictions and parameters on the AI tool's data use and provide documentation of this [14].

Warning: Be mindful of the potential for AI tools to generate indirect identifiers in a dataset.

4.6 Limitations

No matter the type of review (e.g., delegated, full board) conducted for a study, an AI tool should not be given access to any of the following data in an identifiable format, or in combination with indirect identifiers that could potentially identify a participant: biospecimens, such as blood samples; genomic data; personal health information; information or data that could put participants at risk, such as discussions about sensitive or illegal topics. In general, researchers are recommended to limit sharing sensitive, identifying, or private information when using these tools in research with human participants [14]. However, Mi-

Recommendation: AI tool should not be given access to any of these data in an identifiable format.

crosoft Copilot (formerly Bing Chat enterprise) has now been approved for Type 3 data. For more information, see ICT's Copilot webpage.

4.7 Bias

Studies have demonstrated that AI systems can replicate or propagate prejudices depending on a variety of elements, such as the data used to train and manage the AI tool. The REB suggests that researchers create plans to regularly evaluate the AI tools used in their protocols for bias. This testing may include, but is not limited to, constructing test scenarios or replication tests [14].

4.8 Data scraping

The increasing popularity of AI tools for extracting data from websites has been observed. The REB may grant approval for the use of such tools in the context of human subject research, provided the research meets the necessary criteria for approval. Investigators must include the relevant restrictions and parameters of the AI tool in their research proposal [14].

4.9 Interventions

In certain circumstances, the use of AI tools for intervention-based research with participants is permissible, provided the research meets the criteria for approval set by the REB. When submitting to the REB, investigators must include: a comprehensive explanation of the planned interaction between the participant and the AI tool; a description of the data the AI tool will be programmed to collect; documentation of the parameters or limits placed on the AI tool for the intervention, data collection, and (if applicable) data analysis; scripts or texts of instructions that will be read or provided to participants as part of the interaction with the AI tool; and a plan to monitor the safety of participants and their data during and after the intervention. Further, researchers must inform participants of the use and context of AI tools in all consent documents, including any privacy considerations [14].

4.10 Collaborative research

Researchers should take note of this advice when participating in collaborative research that involves AI tools and involves multiple institutions and/or researchers. Laurier's REB will use this guidance to the fullest extent when reviewing the research study. When another REB is responsible for the review, Laurier's REB will use this guidance as part of the "local context review" of the study [14]. Consequently, it is recommended that Laurier researchers ensure that their collaborators at other institutions are aware of Laurier's AI guidance before submitting the study plan to their institutional REBs.

5 Publisher guidance and policy

Advice and regulations from publishers are changing. Before submitting your work, check the particular journal's policy. So far, publishers have been quite clear that generative AI tools should not be used in the editorial process, either by the editor or by reviewers [15, 16, 12, 17]. Publishers have also stated explicitly that generative AI tools should not be listed as authors [16, 18, 17, 19]. Publishers have generally been quiet about the use of generative AI tools to enhance the quality of writing, or have explicitly allowed such use, provided that it is credited properly [15, 12, 17, 19]. However, it must be pointed out that increasingly these tools used to improve writing are themselves using generative AI tools in the background (see also subsection 3.4).[1] Finally, those publishers who have declared a policy on images or videos created by generative AI tools have for now prohibited these uses, with some exceptions [17, 16].

Recommendation: Generative AI tools should not be used in the editorial process, either by the editor or by reviewers, nor should they be listed as coauthors.

[1] Mark Humphries. *Microsoft Copilot Necessitates Some Tough Conversations*. [html]. (Accessed on 02/16/2024)

In a systematic analysis of 2023 of instructions provided by journal publishers to authors on the use of generative AI (GAI) tools, the authors concluded:

> "Among the largest 100 publishers, 17% provided guidance on the use of GAI, of which 12 (70.6%) were among the top 25 publishers. Among the top 100 journals, 70% have provided guidance on GAI. Of those with guidance, 94.1% of publishers and 95.7% of journals prohibited the inclusion of GAI as an author. Four journals (5.7%) explicitly prohibit the use of GAI in the generation of a manuscript, while 3 (17.6%) publishers and 15 (21.4%) journals indicated their guidance exclusively applies to the writing process. When disclosing the use of GAI, 42.8% of publishers and 44.3% of journals included specific disclosure criteria. There was variability in guidance of where to disclose the use of GAI, including in the methods, acknowledgments, cover letter, or a new section. There was also variability in how to access GAI guidance and the linking of journal and publisher instructions to authors. Two journals had GAI guidance that directly conflicted with guidance developed by their publishers." [21]

6 Frequently asked questions

6.1 Can generative AI tools be used in works submitted for publication?

It is important for researchers to be aware of the policies on the use of generative AI in publications. These policies vary among publishers and journals, so it is crucial to follow the particular policies of the pre-print server, journal, or publisher when submitting work (see section 5). Some publishers allow the use of generative AI in the research process, as long as there is a suitable description, references, and supplementary material to show how the AI tool was used. However, they do not allow

AI-generated text to be included. Other publishers allow AI-generated text to be included, but not images. It is generally agreed upon that generative AI tools cannot be considered authors of scholarly works since they cannot take responsibility or be held accountable for submitted work. The Committee on Publication Ethics and the World Association of Medical Editors have issued statements on authorship and the use of AI in scholarly works, which should be consulted. It is important for graduate units, supervisors, and graduate students to understand and comply with the requirements in their field regarding authorship and the use of AI in works submitted for publication or to pre-print servers [22].

Creating text and images for scholarly publications can result in issues of authorship, copyright, and plagiarism, which are not yet resolved. It is a controversial area, and many journals and research conferences are updating their regulations. Therefore, it is essential to read the author guidelines of the journal you are targeting.

Generative AI may serve to facilitate improved writing, especially for non-native English speakers. As long as the human author takes full responsibility for the resulting content, such "editing help" from generative AI is likely to be accepted in most disciplines, where the specific language used is not the focus of the scholarly contribution. However, the use of such techniques may be limited in the short term due to conservative editorial policies at some publication venues. See section 5 for further details.

6.2 What are the copyright and intellectual property implications of using generative AI tools?

The legal situation related to intellectual property and copyright in the context of generative AI is complex and varies across different jurisdictions. It is also evolving rapidly, and the full implications are not yet fully understood. Researchers, including graduate students, must exercise caution while using generative AI tools since certain uses may infringe upon copyright or other intellectual property rights. Moreover, providing data to an AI tool may pose challenges in enforcing intellectual property rights in the future. Generative AI has the potential to create content that plagiarizes the work of others without giving appropriate attribution. Copyleaks reports that nearly 60% of ChatGPT 3.5 generated content includes some form of plagiarism [23]. Graduate students who use AI-generated content in their academic writing may unwittingly include plagiarized material or someone else's intellectual property, which is a violation of Laurier's Policy 12.2 Student Code of Conduct: Academic Misconduct, and Policy 11.14 for the Responsible Conduct of Research. Therefore, it is crucial for them to understand the risks involved [22].

6.3 Who is responsible for AI-generated content used in research or other scholarly work?

Researchers who use AI tools and incorporate the results in their research and written work are ultimately responsible for the content. This responsibility applies to work submitted as part of degree requirements, as well as scholarly publishing or the use of pre-print servers. Authors must be aware of

the terms and conditions of any submission of their work and of any tools they use. These often make the user accountable for the content. This means that researchers may find themselves in a situation where they are accused of spreading false or deceptive information, violating intellectual property rights, violating research ethics approval conditions, committing research misconduct, infringing privacy rights, or engaging in other activities that may carry academic, civil, or criminal penalties [22].

6.4 How do I cite the use of generative AI?

When writing a research paper, it is important to give credit to generative AI for its use. The reader should be informed of its use, and generative AI should not be listed as a co-author. The Committee on Publication Ethics has a clear explanation of this. The paper should include a disclosure of the use of generative AI, along with a description of the places and manners of use. This is usually done in the "Methods" section of the paper. If generative AI output is relied upon, it should be cited, similar to citing a web page or a personal communication. It is important to note that some conversation identifiers may be local to the account, and thus not useful to the reader. The American Psychological Association (APA) and the Chicago Manual of Style have provided good citation style recommendations. [24]

Information: More detailed advice on citation conventions is provided by the Library on their Citing Use of AI (Artificial Intelligence) or Chatbots webpage.

6.5 Can I use generative AI to *write* grants?

An ad hoc panel of external experts advised Canadian federal research funding agencies on the implications of using generative AI in research funding applications [25]. While a blanket ban on the use of generative AI in grant writing is impractical and unenforceable, the panel suggests clear expectations for applicants. These include personal accountability for the complete contents of the application and the development of primary questions and methodologies. The applicant is responsible for framing their work in the broader context of their field and disclosing the use of generative AI in the preparation of a proposal. The panel suggests a simple binary declaration as a starting point: applicants should declare if generative AI was used, for any purpose, in the preparation of the grant. The exact specification of what is to be disclosed requires investigation, but the panel recommends annual review and revision. The panel recommends that Canadian federal research funding agencies adopt a nuanced approach to the use of generative AI in grant writing and evaluation. By doing so, they can protect intellectual property while promoting fairness, transparency, and equity in research funding.

6.6 Can I use generative AI to *review* grant proposals or manuscripts?

In terms of the evaluation of grant applications, the ad hoc panel of external experts advised Canadian federal research funding agencies that funders disallow unsanctioned use of AI in application review to respect intellectual property. Cloud-based generative AI tools necessitate the transmission of the

application's content to the tool and often result in a loss of custody of the intellectual property contained therein. Therefore, reviewers should not use unsanctioned generative AI tools. In-house AI hosting infrastructure is a nontrivial undertaking, but it could be revisited in the future [25].

Internationally, the National Institutes of Health (NIH) has issued a statement prohibiting the use of generative AI to evaluate and review grant proposals. The Government of Canada has followed suit, disallowing the use of generative AI in the grant review process [25]. This applies to both publicly available generative AI systems and those hosted locally, as long as data can be shared with multiple people. The reason behind this prohibition is to maintain confidentiality, which is vital in the grant review process. To use generative AI tools for grant proposal evaluation and summarization or even allowing it to edit critiques, one would need to provide it with "substantial, privileged, and detailed information." However, it is not recommended to do so since we are unsure how the AI system will save, share or use the information that it is fed. Moreover, expert review requires subject matter expertise, which a generative AI system may not possess. Therefore, relying on generative AI to produce a reliable and high-quality review is unlikely. For these reasons, it is not advised to use generative AI for reviewing grant proposals or papers, even if the relevant publication venue or funding agency has not issued explicit guidance [24].

6.7 Can I use generative AI to help me when I write a literature review section for my paper?

Generative AI is a useful tool for researchers as it can quickly summarize a large number of papers and help identify literature related to certain research questions. However, there are a few things to keep in mind when relying on such reviews. Firstly, generative AI may not always have all the necessary information or knowledge to provide an accurate answer, so it's important to verify the summaries it provides. Secondly, unlike human researchers, generative AI cannot evaluate the quality of the published work, so it may include studies of varying quality, some of which may not be reproducible. Thirdly, generative AI has a knowledge cutoff date, so it won't include newer publications after the cutoff date in the responses it gives. Lastly, generative AI's effectiveness depends on the training datasets, and there is no guarantee that the training is unbiased [24].

Uploading full-text articles, book chapters, and other protected content into third-party AI-enabled tools is a challenging issue. There is uncertainty around the exact purposes and uses of storing such content on a third-party server and reusing it through a generative AI tool, as these activities could potentially be infringing copyrights. Most of the big five academic publishers, including Elsevier, Wiley, SpringerNature, Taylor and Francis, and Sage, are explicitly seeking language in their renewal agreements to prevent their content from being used with generative AI and machine language applications. These activities are mostly prohibited by the current licenses that the Library manages on behalf of the University under the broad heading of unauthorized reproduction and third-party licensing.

It is not clear whether uploading full-text content into an AI tool for literature reviews can be considered broadly under the "educational use or purposes" found in most current license agreements, and it carries significant risk. Therefore, we advise researchers not to directly upload content from licensed or restricted electronic resources such as journal articles. In case researchers have any issues related to this matter, they can always request assistance from the Library.

6.8 Can generative AI be used in graduate research?

Graduate students define and investigate an area of research interest to generate new knowledge. Core to this enterprise are values of curiosity, integrity, transparency, and accountability. To be awarded a Master's or Doctoral degree, graduate students are required to demonstrate competencies described by the Ontario Council of Vice-President Academics degree level expectations: *depth and breadth of knowledge, research and scholarship, application of knowledge, professional capacity/autonomy, communication skills and awareness of limits of knowledge*. While generative AI may provide tools to support the attainment of these outcomes, these tools should not replace the ability of graduate students to achieve the degree level expectations.

Recommendation: While generative AI may provide tools to support the achievement of these outcomes, these tools should not replace the ability of graduate students to achieve the degree level expectations.

There are many opportunities for the use of generative AI in graduate research. Investing in (and overcoming) the challenges effective writing presents is a key and central facilitator of the learning process for graduate students conducting research. This includes all stages of the research writing process, from the development of a literature review to reporting the study findings. Although technologies including generative AI can be and are enabling, and should be taken advantage of when deemed appropriate, there is no substitute for the excitement and sense of accomplishment when human ingenuity, curiosity, and critical thinking leads to well-articulated knowledge advancement and innovation. Fostering research skill, both independent of and with the aid of technology, remains integral to the graduate studies experience for research-based graduate programs.

Perception on the use of generative AI varies significantly among individuals and disciplines. Therefore, the Faculty of Graduate and Postdoctoral Studies supports the following broad guidelines defined by responsibilities for Graduate Students, Advisors and Thesis/Dissertation Advisory Committees, and Graduate Programs.

Graduate students have ultimate responsibility for content generated by AI that forms any part of their thesis/dissertation and other degree requirements. Graduate students should be able to replicate and explain all content submitted as part of their academic program. The inability to do so may reflect an academic integrity violation. Graduate students are strongly advised to formally record endorsement from their advisors and committees for the utilization of generative AI in their research and/or writing. If generative AI tools have been used in the research process, there must be a clear statement of which tools have been used, how they have been used, and why they have been used. For example, a declaration should entail: "During the preparation of this work, the author used [NAME TOOL / SERVICE] in order to [REASON]. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the thesis/dissertation."

Recommendation: Graduate students should be able to replicate and explain all content submitted as part of their academic program.

Advisors and advisory committees should provide clear guidance and expectations on what degree of engagement with generative AI in research and/or writing is acceptable (if any). Collectively, the advisor, advisory committee, and graduate student should develop processes by which a student can attest to their own contributions (e.g., during research updates provided at committee meetings), while also noting how the use of generative AI tools contributed to their work.

If a **Graduate Program** permits the use of generative AI in graduate research, discipline-specific norms regarding appropriate use and referencing should be clearly articulated and accessible to advisors, thesis/dissertation advisory committees, and graduate students.

Recommendation: Advisors and advisory committees should provide clear guidance and expectations on what degree of engagement with generative AI in research and/or writing is acceptable (if any). Students considering the use of AI should have a frank conversation about it with their advisor and/or committee *before* starting the work.

6.9 Can I use generative AI to write code?

Generative AI has the ability to output computer programs, but it is crucial to note that the generated code can still contain errors. It is advisable to have a basic understanding of coding before relying entirely on generative AI. However, as reading code is often easier than writing it, generative AI can be a useful tool for writing code for you. This can be applied not only to computer programs but also to databases. By using generative AI to write SQL code, you can easily manage and query databases. A simple way to test the code generated is by running it on known instances and checking the output. However, it's important to keep in mind that even if the code works perfectly on simple cases, it doesn't necessarily mean it will work on complex ones [24].

6.10 Can I use generative AI to write non-technical summaries, create presentations, and translate my work?

Generative AI can be advantageous for summarizing or translating your work, particularly with its capacity to modify the tone of a text, making it simpler to create concise yet comprehensive summaries that are suitable for different types of readers. There are several advanced generative AI models that are specifically designed to convert scientific manuscripts into presentations. However, it is important to ensure that confidential information is not inputted into generative AI when summarizing, presenting, or translating your work. Additionally, it is essential to verify that summaries, presentations, and translations created by generative AI accurately reflect your work. When using generative AI for translation, it can be difficult if you are not proficient in both languages involved and you need to consult with a fluent speaker for verification. Additionally, not all generative AI models are explicitly designed for translation tasks, so it is important to explore and identify the most suitable generative AI model that meets your specific translation needs [24].

6.11 How secure are generative AI tools?

Using third-party generative AI tools to process sensitive data may pose privacy and security risks. Researchers working with such data should avoid submitting personal or identifying participant information to these tools as it could breach research participant privacy. Similarly, confidential information should not be shared with these tools without documented permission. Researchers should seek assistance in assessing the risk associated with using third-party AI tools and can find resources at ICT, the Library, or the Centre for Teaching and Learning [22] .

6.12 How do I decide which generative AI to use in research?

The most important factor to consider when selecting a generative AI system is which one best fits your research questions (what data, what model, what computing requirements). Additionally, there are some general considerations to take into account. *Open source*. Open source software is published with its source code, allowing anyone to use and explore it. This is important for generative AI models, as they are usually not developed locally by the researchers themselves (unlike Machine Learning models). Open source generative AIs, as well as generative AI systems trained with publicly accessible data, can be beneficial for researchers who want to fine tune the models, assess the security and functionality of the system, and improve the explainability and interpretability of the models. *Accuracy and precision*. When the outputs of a generative AI can be verified (e.g. for data analytics), its precision and accuracy can be evaluated. *Cost*. Some models require subscriptions to APIs for research use, while others may be integrated locally but come with integration and maintenance costs. When selecting free models, you may need to cover the cost of an expert to set up and maintain the model [24].

It is important that researchers follow the Privacy and Security Impact Assessment (PSIA) Process as detailed in ICT's Finding Software Solutions webpage.

6.13 What uniquely generative AI issues should I consider when I use it in my research?

The implications of generative AI are numerous and require careful consideration. *Ethical issues* such as data privacy, transparency, and accountability are of particular concern, as generative AI is often used in closed-source systems. Additionally, *bias in the data* and the output of the AI system can be a major issue, as generative AI is trained on large datasets that are often inaccessible and may contain biases and stereotypes. Furthermore, most generative AI models are *trained with predominantly English texts, western images, and other types of data*, which means that non-Western or non-English speaking cultures, as well as work by equity-deserving groups and non-English speakers, are underrepresented. *AI hallucination* is another issue, as generative AI can produce outputs that are factually inaccurate or entirely incorrect, uncorroborated, nonsensical or fabricated. *Plagiarism* is also a potential problem, as generative AI can only generate new content based on or drawn from the data on which it is trained. *Prompt engineering* is a new human activity that has arisen due to generative AI, as the quality of generative AI responses is heavily influenced by user input or 'prompt.' Furthermore, generative AI models are *trained on data up to a specific date*, meaning that they are unaware of any

events or information produced beyond that date. *Model continuity* is also a factor to consider, as external entities or vendors can discontinue the model, which affects the reproducibility of research. Finally, *security* is a major concern, as generative AI systems are susceptible to security breaches and attacks, such as prompt injection, which deliberately feeds harmful or malicious content into the system to manipulate the results [24].

7 Additional resources

7.1 Journal articles & reports

Council of Canadian Academies, 2022. *Leaps and Boundaries*, Ottawa (ON). The Expert Panel on Artificial Intelligence for Science and Engineering, Council of Canadian Academies. [pdf]

Buriak, Jillian M., Deji Akinwande, Natalie Artzi, C. Jeffrey Brinker, Cynthia Burrows, Warren C. W. Chan, Chunying Chen, et al. "Best Practices for Using AI When Writing Scientific Manuscripts." *ACS Nano* 17, no. 5 (2023): 4091–93. [doi].

Flanagin, Annette, Kirsten Bibbins-Domingo, Michael Berkwits, and Stacy L. Christiansen. "Nonhuman 'Authors' and Implications for the Integrity of Scientific Publication and Medical Knowledge." *JAMA* 329, no. 8 (February 28, 2023): 637–39. [doi].

Foltynek, T., Bjelobaba, S., Glendinning, I. et al. ENAI Recommendations on the ethical use of Artificial Intelligence in Education. Int J Educ Integr 19, 12 (2023). [doi]

Harker, Jennifer. "Science Journals Set New Authorship Guidelines for AI-Generated Text." *Environmental Factor* (blog), November 15, 2023. [html].

Hosseini M, Horbach SPJM. Fighting reviewer fatigue or amplifying bias? Considerations and recommendations for use of ChatGPT and other large language models in scholarly peer review. Res Integr Peer Rev. 2023 May 18;8(1):4. [doi]. Erratum in: *Res Integr Peer Rev*. 2023 Jul 10;8(1):7. PMID: 37198671; PMCID: PMC10191680.

Hutson, Matthew. Could AI help you to write your next paper? *Nature* 611.7934 (2022): 192-193.

7.2 Introduction to generative AI

The following reading list is recommended by Google:

- Ask a Techspert: What is generative AI? [html]
- Build new generative AI powered search & conversational experiences with Gen App Builder: [html]
- What is generative AI? [html]
- Google Research, 2022 & beyond: Generative models: [html]
- Building the most open and innovative AI ecosystem: https://cloud.google.com/blog/products/ai-machine-learning/building-an-open-generativ e-ai-partner-ecosystem
- Generative AI is here. Who Should Control It? [html]

- Stanford U & Google's Generative Agents Produce Believable Proxies of Human Behaviors: [html]
- Generative AI: Perspectives from Stanford HAI: [html]
- Generative AI at Work: [pdf]
- The future of generative AI is niche, not generalized: [html]
- The implications of Generative AI for businesses: [html]
- Proactive Risk Management in Generative AI: [html]
- How Generative AI Is Changing Creative Work: [html]

7.3 Large language models

The following reading list is recommended by Google:

- NLP's ImageNet moment has arrived: [html]
- LaMDA: our breakthrough conversation technology: [html]
- Language Models are Few-Shot Learners: [pdf]
- PaLM-E: An embodied multimodal language model: [html]
- PaLM API & MakerSuite: an approachable way to start prototyping and building generative AI applications: [html]
- The Power of Scale for Parameter-Efficient Prompt Tuning: [pdf]
- Google Research, 2022 & beyond: Language models: [html]
- Solving a machine-learning mystery: [html]

7.4 Miscellaneous

- Attention is All You Need: [html]
- Transformer: A Novel Neural Network Architecture for Language Understanding: [html]
- Transformer on Wikipedia: [html]
- What is Temperature in NLP? [html]
- Model Garden: [html]
- Auto-generated Summaries in Google Docs: [html]

Reference List

[1] Kim Martineau. *What is generative AI? | IBM Research Blog*. [html]. (Accessed on 02/19/2024).

[2] *Generative AI Usage Guidance for the Research Community - Provost & Chief Academic Officer - UNC Chapel Hill*. [html]. (Accessed on 11/15/2023).

[3] Beata M Jones. *How Generative AI Tools Help Transform Academic Research | Forbes Magazine*. [html]. (Accessed on 03/26/2024). Sept. 2023.

[4] Auro Del Giglio and Mateus Uerlei Pereira da Costa. “The use of artificial intelligence to improve the scientific writing of non-native english speakers”. en. In: *Journal of The Brazilian Medical Association* 69.9 (Sept. 2023), e20230560.

[5] Yonah Welker. *Generative AI holds great potential for those with disabilities — but it needs policy to shape it | World Economic Forum*. [html]. (Accessed on 03/26/2024). Nov. 2023.

[6] Government of Canada. *Guide on the use of Generative AI - Canada.ca*. [html]. (Accessed on 02/16/2024).

[7] Dan O’Keefe. *Private AI vs. Public AI: 4 Things to Know*. [html]. (Accessed on 02/19/2024).

[8] *Why You Shouldn’t Trust ChatGPT With Confidential Information*. [html]. (Accessed on 11/18/2023).

[9] Garling Wu. *8 Big Problems With OpenAI’s ChatGPT*. [html]. (Accessed on 11/17/2023). Sept. 2023.

[10] *Introducing Microsoft 365 Copilot – your copilot for work - The Official Microsoft Blog*. [html]. (Accessed on 11/24/2023).

[11] Association for Computing Machinery. *Frequently Asked Questions*. [html]. (Accessed on 02/26/2024).

[12] *ChatGPT and Generative AI | SAGE Publications Inc*. [html]. (Accessed on 11/19/2023).

[13] Sunil Ramlochan. *The Truth About AI Detectors — More Harm Than Good*. [html]. (Accessed on 03/25/2024). 2023.

[14] *Artificial Intelligence (AI) Tools - Research Integrity & Assurance*. [html]. (Accessed on 11/17/2023).

[15] *Publishing ethics | Elsevier policy*. [html]. (Accessed on 11/19/2023).

[16] *Artificial Intelligence (AI) | Nature Portfolio*. [html]. (Accessed on 11/19/2023).

[17] *Science Journals: Editorial Policies | Science | AAAS*. [html]. (Accessed on 11/19/2023).

[18] *Defining authorship in your research paper - Author Services*. [html]. (Accessed on 11/19/2023).

[19] *Best Practice Guidelines on Publishing Ethics | Wiley*. [html]. (Accessed on 11/19/2023).

[20] Mark Humphries. *Microsoft Copilot Necessitates Some Tough Conversations*. [html]. (Accessed on 02/16/2024).

[21] Conner Ganjavi et al. *Bibliometric Analysis of Publisher and Journal Instructions to Authors on Generative-AI in Academic and Scientific Publishing*. 2023. arXiv: 2307.11918 [cs.DL].

[22] *Guidance on the Appropriate Use of Generative Artificial Intelligence in Graduate Theses – School of Graduate Studies*. [html]. (Accessed on 11/15/2023).

[23] *Copyleaks Research Finds Nearly 60% of GPT-3.5 Outputs Contained Some Form of Plagiarized Content | Press Release*. [html]. (Accessed on 02/27/2024).

[24] *A Quick Guide of Using GAI for Scientific Research | MIDAS*. [html]. (Accessed on 11/15/2023).

[25] *Advice from the ad hoc generative AI panel of external experts*. [html]. (Accessed on 02/19/2024).

Index