

Ursinus College Responsible AI Principles

Ursinus College remains committed to centering the human in education. Deep, individual thought and person-to-person connection and collaboration continue to be at the forefront of an Ursinus education, and in fact, in the face of new technologies, are more critical than ever.

New technologies, including generative artificial intelligence (AI) and large language models (LLMs), are becoming increasingly interwoven into our society. We therefore seek to equip a student body, faculty, and staff to use these tools capably, ethically, and mindfully, in line with [Ursinus's mission](#) to “enable students to become independent, responsible, and thoughtful individuals through a program of liberal education.”

To that end, the following document offers definitions of AI and LLMs; considerations for maintaining privacy; suggestions for using AI responsibly if you involve it in your teaching and students’ learning; samples of policies you can adopt or adapt, depending on how much AI you want to permit students to consult; and guidelines for citing and detecting AI.

The college encourages students, faculty, and staff to remain mindful of privacy, equity, security, and ethics when exploring AI and LLMs, and to ensure that each decision about its use remains compatible with the college’s mission.

Contents

Definitions	2
College-Wide Policy for AI Use in Student Work.....	2
Minimizing Privacy Risks with Generative AI.....	3
Responsible AI Use in Teaching and Learning	4
General Considerations	4
Academic Integrity and Clear Communication of Expectations.....	4
Example Course Policy Language for Generative AI Use	5
Disclosing AI Usage	6
Detecting AI-Generated Content	6
Guidance for Students.....	7
Further Support.....	7
Work Consulted	7

Definitions

Generative Artificial Intelligence (AI) is an emerging and rapidly evolving type of artificial intelligence that uses machine learning models to create transformed or novel content, such as text, images, music, or videos, based on user prompts or inputs. A **Large Language Model (LLM)** is a type of generative AI that is trained on vast amounts of data to enable it to predict a probable sequence of words in response to a prompt and produce coherent, contextually relevant responses. **Agentic AI** uses specialized agents to perform specific tasks across multiple software platforms and can operate like a group of collaborating workers. For the purposes of this guide, uses of the term “AI” will refer to generative AI, like ChatGPT or Claude.

College-Wide Policy for AI Use in Student Work

At Ursinus, the decision whether to permit students to use AI tools to complete coursework or assignments resides with the professor. Professors should state their AI policy on the syllabus, and students should find out from the professor which assignments and tasks allow the use of AI and which do not. Students who use AI with the permission of their professor should disclose and cite its usage according to the guidance of a citation style guide, such as MLA, APA, or Chicago, or as the professor

instructs. Cases of academic dishonesty involving the use of AI will be treated as any other academic dishonesty case, following the college's handbook-specified [procedures for suspected academic honesty violations](#).

Minimizing Privacy Risks with Generative AI

The Office of the Provost, the Teaching and Learning Institute, and Information Technology recommend that faculty, students, and staff proceed with caution when using generative AI tools such as OpenAI's ChatGPT, Google's Gemini, Adobe's Creative Cloud, Anthropic's Claude, and Microsoft's Co-Pilot, as well as others. The guidance offered below focuses largely on mitigating risks, some of which carry serious social, academic, and/or legal consequences.

Many third-party software tools do not offer robust privacy controls. If using public AI chatbots, please follow the guidelines below:

- In interactions with these tools, do not use or provide any information that you should not share publicly. Examples include financial statements, payroll or salary details, disciplinary records, college processes, and system architecture and/or vulnerabilities. Faculty and staff should not require students to use AI tools in any way that requires them to reveal such information about themselves or the college.
- Do not use personally identifiable information or Ursinus' institutional data, such as donor history, student transcripts, tenure review details, resumes or application letters, payment history, etc. Faculty and staff should not require students to use AI tools in any way that requires them to reveal such information about themselves or the college.
- If you are considering licensing a new tool or using a new AI feature in a tool you already use, please work with [IT](#) to ensure that the tools and services you procure or are using on behalf of the College have appropriate privacy and security protections and are assessed for risk prior to use.

Ursinus College's licensing agreements with OpenAI and Anthropic ensure that prompts and files submitted to those companies through [Ursinus.ai](#) are not kept by those companies, nor are they used to train AI algorithms. All documents and prompts stored on [Ursinus.ai](#) remain on college-controlled servers and are accessible to the college but cannot or will not be shared with or used by third parties. Faculty and staff with questions about how Ursinus servers store or view [Ursinus.ai](#) prompts can refer to the college's IT [privacy policy](#); reach out to IT directly with further questions.

Ursinus IT currently recommends *against* using agentic AI for college-related tasks, since AI agents require access to password-protected information systems to complete

their tasks autonomously.

Responsible AI Use in Teaching and Learning

General Considerations

If using third-party generative AI tools or [Ursinus.ai](#) as part of teaching and learning, faculty should consider the following:

- Artificial intelligence tools can be inaccurate. They can produce misleading or completely fabricated information that was not in the input data. Furthermore, they may exhibit biases.
- Always be sure to verify the suggestions or information AI tools provide.
- Be sure to comply with Ursinus' academic integrity policies.
- [Early studies](#) have suggested an overreliance on AI tools can impede learning and critical thinking. Take the time to consider how your use of AI will supplement and enhance but not undermine students' learning process and critical thinking.
- If you have already begun the use of new tools or features, configure the platform *not* to use or retain chat/prompt history, if possible. Once data are placed into these platforms, there are no guarantees that it can later be removed from them.
- Consider modeling responsible AI use for your students by disclosing when you have used AI to generate course content, activities, or other materials that they interact with.

Academic Integrity and Clear Communication of Expectations

The generative AI landscape is evolving rapidly and is increasingly being integrated into the academic lives of both students and faculty. As a result, all instructors should provide clear guidelines concerning the use of generative AI in academic work for every class they teach. Here are a few things to keep in mind when addressing AI and academic integrity in your course.

To reduce the chance of violations of academic integrity, explicitly communicate your course's generative AI policies by taking the following steps:

- Include clear guidelines in your syllabus regarding the use of generative AI tools, specifying any differences in the usage policy for particular assignments. Discuss the expectations outlined in the syllabus during a class session.
- Review expectations and communicate how they relate to specific assignments. Remember that because students have many classes with differing AI policies, they're likely to need multiple reminders of your policy.
- Explicitly state when the use of generative AI is allowed or prohibited.

- If you are concerned about the intellectual property concerns that arise if students share your course materials with an AI product, specify any prohibitions in your syllabus and re-iterate them throughout the semester.
- When generative AI is allowed, clarify expectations for student documentation and attribution, specifying which parts of the work must be done by the student without AI assistance and explaining how they should validate or verify the permitted AI-generated content.
- If needed, engage in ongoing discussions about the importance of academic integrity, emphasizing that core academic integrity principles remain crucial and apply even with the availability of generative AI tools.

Example Course Policy Language for Generative AI Use

When planning your syllabus and course policies regarding generative AI, consider adapting the following language to convey a general stance in your syllabus. Keep in mind that this sample language provides a course-level overview on broadly allowing or disallowing the use of generative AI tools.

Prohibiting Generative AI Use in Your Course	Example language
	To ensure that you develop and master the foundational concepts and skills in this course, the use of generative artificial intelligence (AI) tools is not allowed. If you are uncertain about any policy or specific assignment directions—including whether a tool qualifies as generative AI (like Grammarly or AI summaries from search engines)—please consult me before using the technology or completing your assignment.
Permitting Generative AI Use with Attribution in Your Course	Example language
	Mastering the essential foundational concepts of this course requires effort and practice. Therefore, the use of generative artificial intelligence (AI) tools is prohibited in most cases, but may be allowed for select activities or assignments. Whether generative AI assistance is permitted for a specific assignment will be clearly communicated when that assignment is introduced. You should cite any uses of generative AI according to the style guidelines provided by [specify citation style here].

Encouraging Generative AI Use with Attribution in Your Course	Example language
	The use of generative artificial intelligence (AI) tools is encouraged for specific activities or assignments. Each assignment's instructions will clearly state whether or not generative AI is allowed. You should cite any uses of generative AI according to the style guidelines provided by [specify citation style here].

Disclosing AI Usage

If you use material generated with AI in course materials, model the behavior you expect from your students. If you expect transparency from your students, then your transparency is key: each instance should be properly referenced.

If you use AI to develop course materials, consider doing the following:

- 1) Explain at the beginning of the document how you used AI in the development of the course material.
 - E.g. In creating this handout, I used Anthropic's Claude to generate the examples of "bad cover letters," rather than using real student-generated work. Each place where I included AI-generated text, I added a footnote that indicates so.
- 2) Cite the AI-generated portions, following the conventions of your field's citation guidelines. You may want to select citation guidelines appropriate to your discipline and share them with your students. Currently, the [APA style guide](#), [MLA style guide](#), and [Chicago Manual of Style](#) include recommendations and examples for citing AI-generated materials. For example, Chicago style recommends stating in a footnote when text was generated by an AI tool (e.g., Footnote: Text generated by ChatGPT, OpenAI, and edited by the author for style and clarity).
- 3) In short, in your own work, strive to meet the same expectations you have for your students.

Detecting AI-Generated Content

It can be quite difficult to clearly detect AI-generated content using detection technology. Many third-party detection tools claim to identify work as AI-generated, but cannot provide evidence for that claim, and follow-up assessments have revealed significant margins of error. This raises a substantial risk that students may be wrongly accused of improper use of AI. Therefore, given their unreliability and current inability to provide definitive evidence of violations, we do not recommend using automatic detection algorithms for academic integrity violations involving AI use.

Instead, we urge you to establish trusting relationships with students and emphasize learning as a process based on relationships and effort. Consider designing [authentic assessments](#), valuing process over product, introducing [AI transparency](#) into your assignments, and establishing community agreements or [AI contracts](#) for AI tool use and behavior in the course, which can be more effective than policing students.

Professors can still sometimes detect dishonest use of AI, either because the submitted work seems out of step with the student's other work or because it exhibits traditional traits of plagiarized work. If you think a student used AI in a way that violates the assignment's policies or academic integrity, follow the same process you would for any violation of academic integrity, according to the policy laid out in your syllabus.

Guidance for Students

A version of this document has been prepared for students and can be accessed on the Canvas AI Resources page.

The Center for Writing and Speaking has a guide for citing generative AI that can be accessed [at this webpage](#), under "Documentation."

Further Support

You can refer to the Ursinus College [AI Faculty Resource Canvas page](#) for ideas on ways to incorporate critical AI literacy instruction into your course assignments, as well as AI best practices, colleagues' assignments that incorporate AI, a discussion board for questions, and other related resources.

If you have questions about the technical aspects of AI or any other questions, please contact IT at techsupport@ursinus.edu. If you have questions about the pedagogical aspects of AI, please contact the director of the Center for Writing and Speaking at targondezzi@ursinus.edu, the instructional technology librarian at mgood@ursinus.edu, or the directors of the Teaching and Learning Institute at tlirectors@ursinus.edu.

Work Consulted

Portions of this guide were adapted, with permission, from Colby College's "Working Paper: Guidelines for Principles, Definitions, and Guidance For Working With Generative Artificial Intelligence." Most sections have been modified extensively from the original to reflect the Ursinus context.