

Syllabus Statements About Appropriate AI Use - Guidance for Instructors

All course syllabi at both undergraduate and graduate levels should include information about permissible and appropriate AI use.

A strong AI course policy statement contains the following elements:

1. **Scope and Definition:** what types of AI tools does the policy apply to?
2. **Rationale:** what is the intention behind the policy?
3. **Permissions:** what specifically is and is not allowed?
4. **Verification and Responsibility:** what do you expect in terms of students validating AI outputs?
5. **Disclosures and Citations:** how do you want students to document AI use?
6. **Seeking Advice:** a reminder that students should reach out with questions
7. **Integrity and Consequences:** a reminder that failure to follow the course policy constitutes an academic integrity violation

Below we provide examples for each element, which can be mixed, matched, and revised to suit your course-specific goals and intentions. Three examples of complete syllabus statements (for an AI-forward course, a bounded-use course, and a no-AI course) can be [found here](#).

A) Scope and Definition

As AI tools become ubiquitous, it is important to define for students what your course's AI policy applies to. The variants below provide examples of different scope.

Functional variant

In this course, "AI" means any tool that does substantial work you would otherwise have done yourself — whether by generating content (writing, code, images, analysis) or by carrying out the steps of a task on your behalf (an "agent" that browses, runs code, or works through a problem). The brand does not matter and the list changes constantly; the test is what the tool does for you, not what it is called. [Name any specific tools you want to place clearly on one side of that test.]

Boundary-forward variant: what counts

"AI" is easy to spot in its obvious forms, including ChatGPT, Claude, Gemini, and similar tools that write or create on request. The harder questions arise now that AI is built into everyday software: grammatical suggestions, predictive text, search summaries (e.g., at the top of a Google search), translation, coding autocomplete. A newer hard case is agentic AI — tools that don't just answer but go and do: browsing, running code, completing multi-step tasks. For this course, the rule is: [e.g., "a tool counts as AI use if it drafts or rewrites

substantive content for you; handing an assignment to an agent to carry out counts as AI use here, however much or little it returns to you."] When you are unsure which side a tool falls on, ask before using it.

Boundary-forward variant: what doesn't count

What does not count as AI use for this course includes simple and hard-to-avoid integrations into everyday software such as: [grammatical suggestions, predictive text, search summaries (e.g., at the top of a Google search), translation, coding autocomplete].

Broad variant

In this course, "AI" is defined broadly: any tool that generates, rewrites, translates, paraphrases, substantially reworks content, or carries out tasks for you in response to a prompt. This includes chatbots such as ChatGPT, Claude, and Gemini, as well as AI paraphrasing and translation tools. It also includes agentic tools that complete multi-step work on your behalf. If a tool produces language, ideas, or other outputs you would otherwise have produced yourself, treat it as covered.

Narrow variant

In this course, "AI" means standalone tools that generate original content from a prompt — for example ChatGPT, Claude, and Gemini. It also means tools that carry out multi-step tasks on your behalf ("AI agents"). Ordinary assistive features such as spell-check, grammar suggestions, autocomplete, and search summaries (e.g., at the top of a Google search) are not included unless noted for a specific assignment. [Add any tool you want to treat differently.]

B) Rationale

Students are more likely to respect and follow a policy when they understand the underlying reasons you are adopting it. Your rationale should link to the specific learning objectives and prohibitions of your course.

Embrace variant

This course treats AI models as tools worth learning to use well. Knowing when these tools genuinely help, when they get in the way, and how to check what they produce is itself a skill this course is meant to build — so you are encouraged to use them within the limits below.

Limited-use variant

Permissible use of AI in this course follows a key principle: the core intellectual work must be performed by you and you alone because it is only by doing this work that you will develop the skills this course is designed to develop. AI is welcome for the supporting tasks named below and off-limits for core intellectual work.

Prohibit variant

Use of AI tools is not permitted in this course because the work these tools would do for you is the very work the course exists to help you learn to do yourself. Outsourcing this work would mean leaving without the skills you came to build.

C) Permissions

When some but not all AI use is permitted, be as specific and concrete as possible about the conditions under which it is and is not allowed.

Matrix form

Assignment or task	AI	Conditions
First drafts of papers	Not allowed	—
“Peer-review” on written drafts	Allowed	Feedback from AI tools must be submitted with the final paper, along with commentary about how you responded to the feedback.
Exams	Not allowed	—

Prose form

You may use AI tools for [tasks / assignments]. You may not use them for [tasks / assignments], which include the core intellectual work of this course — [e.g., developing a thesis, gathering and weighing evidence, structuring your argument]. When in doubt, treat AI use for a task as not permitted until you have checked with me.

More [examples available here](#).

D) Verification and Responsibility

When AI use is permitted, it is advisable to include information about students’ responsibility to verify their output and ensure accuracy.

These tools can produce fluent, confident outputs that are wrong, including invented facts, fake quotations, and citations to sources that do not exist. You are responsible for everything you submit, including any errors or fabrications a tool introduces. Check every AI-derived claim against a reliable source before you rely on it.

If you let an agentic AI tool carry out steps you did not directly oversee (e.g., sources it read, code it ran, choices it made along the way), you remain responsible for all of it and you must be able to explain what it did.

E) Disclosure and Citation

If you want students to disclose their use of AI tools, provide information about how they should do so. Disclosures can range from minimal to extensive.

Simple Disclosure

When you use a permitted tool, disclose it. One line is enough: name the tool, say what you used it for, and say what you did with the output. For example: *"I used ChatGPT to brainstorm counterarguments, then chose and developed my own."*

Process Disclosure

Append a paragraph at the end of the assignment describing how AI was used across the assignment, and how you checked it. For example: *"I used Gemini to summarize three sources and to flag errors in my code. I verified the summaries against the originals and rewrote them in my own words, and I tested and corrected the code myself."*

Prompts Disclosure

Attach the actual prompts you used as an appendix.

Delegation Disclosure

When you delegate a task to an agent, a one-line note about the tool isn't enough. Describe the task you delegated, how much you let it do on its own, and what you did or didn't verify. For example: *"I used an agent to gather and screen sources; I let it run the searches and first-pass filtering, then read and re-checked the dozen it kept. I did not independently verify the ones it discarded."*

Citation

If you quote or reproduce AI-generated content, also cite it formally (see [How to Cite ChatGPT & Similar AI Tools](#)). If you are unsure whether something needs citing, cite it anyway.

F) Seeking Advice

Encourage students to reach out to you when they are unsure about what is and is not appropriate or permissible.

If you are not sure whether a particular use is allowed, ask me before you do it. I would far rather answer the question in advance than have to treat an honest mistake as a violation.

G) Integrity and Consequences

Be explicit about the fact that not following the course AI policy constitutes an academic integrity violation and the consequences that will follow.

Using AI tools in ways this course does not permit is a violation of [Lehigh's Code of Conduct for Academic Integrity](#), which prohibits, among other things, using unauthorized tools or resources; using or paraphrasing the work, themes, or ideas of another without full and clear acknowledgment; and submitting the work of another as your own. All violations of Academic Integrity will be referred to the Office of Student Conduct and Community Expectations and will result in disciplinary action.