

AI POLICY

For certain essay assignments indicated in the syllabus, you may use large language models (LLMs) such as ChatGPT, Claude, NotebookLM, etc., but only under the constraints that you will find below. Their aim is to preserve the core intellectual work of doing philosophy—forming a view, building an argument, and defending it—while employing LLMs to function as tools for refinement and critique. Following this LLM policy falls under the College’s honor code.

1. Your ideas and argument structure must come first.

You must do the initial intellectual work on your own, without LLM assistance. Before using any LLM for this assignment you should already have:

A thesis, i.e., what you are arguing for or against, and
A working outline of the paper, namely the main sections, the path of the argument, and at least a sketch of one or two objections you plan to address.

You may not ask an LLM to “come up with a topic,” “generate a thesis,” or “outline my paper for me.” The central argumentative architecture must originate in your own thinking.

A good internal check: you should be able to explain your project, out loud and without notes, before you ever open an AI tool.

2. No cutting and pasting.

Anything that appears in your paper must be typed by you into your document. You may not cut and paste text from any AI tool directly into your draft.

If you get a helpful objection, example, or formulation from an LLM, you must:

Digest and understand it, and
Rewrite it in your own words, in light of your understanding and the rest of your project.

If you cannot explain an idea in your own words without looking back at the AI output, you are not yet entitled to use that idea in your paper.

3. Permitted uses of LLMs (after you have a thesis and outline)

Refining expression and organization.

Getting help tightening prose, clarifying sentences, improving transitions, or reorganizing paragraphs for clarity.

Asking the model to propose possible objections, counterexamples, or alternative perspectives on your thesis and argument. You must determine which ones are interesting, pressing, or worth taking up.

Summarizing material you have already read.

For example, you may use NotebookLM to get a compressed reminder of a long article or to see a rough map of its main claims.

These summaries are aids to your own understanding, not replacements for doing the reading.

Any characterization of an author's view that appears in your paper must ultimately be grounded in your own engagement with the text and checked against it.

Limited conceptual clarification.

You may ask for rough, introductory explanations of background concepts or distinctions, with the understanding that these are not authoritative sources. Any view or text you rely on must be checked against the primary or secondary literature and cited accordingly.

In all of these cases, you must remain the thinker in charge. You decide what is helpful, what is confused, and what belongs in your paper.

4. Prohibited uses

No AI-written replies to objections.

You may use an LLM to generate objections, but the responses to those objections must be entirely your own work. Do not ask an LLM to "write a reply to this objection" and then use that reply, even if you rewrite or polish it.

No ghost-writing of content. You may not:

Ask an LLM to “write a paragraph (or section) defending my thesis” and then submit that as part of your paper, even in altered form.

Ask an LLM to “rewrite this section to be more persuasive” where what is really changing is the substantive line of thought, not just phrasing.

No AI as a cited scholarly source.

You may not cite an LLM as an authority on a scholar’s view, a text, or a historical claim. All substantive claims in your paper must ultimately trace back to readings, lectures, or scholarship that you have actually consulted.

5. Transparency and AI use statement

You must include, as an appendix at the end of your paper, a brief “AI Use Statement” (a short paragraph or bullet list) that:

States which tool(s) you used (e.g., “ChatGPT,” “Claude,” “NotebookLM”).

Describes what stages of the process you used them for (e.g., “brainstorming possible objections,” “polishing prose in section 3,” “summarizing an article by X that I had already read,” “asking for a simpler explanation of Kripke’s notion of rigidity”).

Includes sample prompts if your use was substantial (for example, you may copy a couple of representative interactions).

If you used AI only in a trivial way (e.g., checking spelling and grammar), you do not need to include an AI use statement.

6. A practical test

As you work, keep asking yourself:

Did I have a clear, reasonably complete understanding of my thesis and main argument before I opened an LLM?

Am I using the model as a tool to help me express and interrogate my own thinking, or as a shortcut to avoid doing that thinking myself?

Could I produce what is in this paper—ideas, structure, and prose—without AI, even if it would take longer?

If you're unsure whether some use is appropriate under this policy, the safe options are: (i) don't do it, or (ii) ask me.