

Richard Williams's Course AI Policy

Last revised August 22, 2025

Prepared by Richard Williams, Notre Dame Sociology, rwilliam@nd.edu

Web page: <https://academicweb.nd.edu/~rwilliam/>

Policy is subject to revision in the future, maybe even during the semester. Comments are welcome.

An annotated bibliography documenting the claims made about AI is available at <https://academicweb.nd.edu/~rwilliam/AIConcerns/AIBibliography.pdf>.

I know this section is long but I want you to read it carefully so you will be sure to understand it. Further I will ask you questions about AI and your comments may help me make improvements.

When used properly, Artificial Intelligence (AI) can be a powerful tool. AI can quickly answer basic questions, synthesizing key information from multiple sources. It often provides links to those sources, allowing you to investigate further. Once you have a draft of your answer, you can use AI (which includes programs like Grammarly) to refine and improve your writing.

When used improperly, however, AI can have very negative effects in a classroom. It can lead to superficial learning or no learning at all. It can lead to students making terrible and embarrassing mistakes. Even the environment can suffer from excessive use of AI.

Before using AI, students should understand the following.

I. When it comes to learning, research shows that AI can be inferior to more traditional approaches. A 2024 review of the literature found that “Over-reliance on AI can lead to diminished creativity and critical thinking abilities, as students may become too dependent on AI-generated content and less engaged in developing their ideas... This dependency can foster complacency and reduce essential problem-solving skills.”

An experiment conducted by researchers at MIT made the potential harms of AI even clearer. In the experiment, some randomly chosen students used AI to write essays, some used search engines, while still others used neither (i.e., brain-only). The results were disturbing:

- “The most consistent and significant behavioral divergence between the groups was observed in the ability to quote one's own essay. LLM users significantly underperformed in this domain, with 83% of participants (15/18) reporting difficulty quoting in Session 1, and none providing correct quotes... Search Engine and Brain-only participants did not display such impairments... *If users rely heavily on AI tools, they may achieve superficial fluency but fail to internalize the knowledge or feel a sense of ownership over it.*”
- “ChatGPT users showed dramatically weaker brain connectivity... and when they later tried to write without AI, their brains looked more like novices than practiced writers. Like steroids that make your hit strong, but your arms flabby.”

I do not find these results surprising. With modest user effort, AI may be able to quickly crank out a nice paper, but how much of it are you going to remember a day later?

I'm sure the harms to learning are even greater when students let AI do all the work for them. Sure, AI makes things easier, but so would hiring somebody to write your paper for you! In classes like mine, which emphasize discussion and students learning from each other, some people using AI to make things as easy as possible for themselves can potentially damage everyone's learning. Employers may want you to know how to use AI, but they probably also want you to show that you actually learned something in all those classes somebody spent a fortune on for you to take.

II. AI frequently hallucinates!!! Just because it doesn't know the answer doesn't mean it won't give you one! There are numerous horror stories that illustrate this.

- A student of mine had an interesting, plausible-sounding, seemingly well-written discussion of one of the class readings. Unfortunately for the student, I had actually read the article I'd assigned, and recognized immediately that the answer was total nonsense. If the student had read even the first few sentences of the article they too would have realized the answer was absurd, but apparently that was just too much darn hard work. The student failed the final exam and ended up barely holding on to a C for the course.
- The Chicago Sun Times recently published a recommended summer reading list, but it didn't turn out so well. "So... the Chicago Sun-Times, which recently laid off 20% of its staff, has published an article titled Summer reading list for 2025. It lists books by well-known authors. There's only one problem: while the authors are real, most of the books are not. Clearly, with the absence of journalistic staff, they resorted to the short cut of asking ChatGPT — and ChatGPT did what it does best: it hallucinated a bunch of nonexistent citations."
- These incidents are embarrassing, maybe even a little amusing. But the costs of AI mistakes can be far greater. A researcher on climate security had an even worse story to tell. "As a researcher active in the climate security field, I recently conducted my own experiments with GenAI. The results were not merely disappointing; they were alarming. Despite the existence of easily accessible information, both GenAI platforms I used drew upon fictional data – and did not make this usage transparent..."

"As someone who has worked in the field for many years, I picked up these issues with GenAI results quickly. But students, researchers, and decision makers less familiar with climate security research would be likely to take such responses at face value. The replies provided by Copilot and ChatGPT sounded plausible, confident, and nuanced, and neither AI platform admitted any error until I explicitly pointed to them. In a field of knowledge like climate security, in which the stakes are high, this suggests that GenAI should never be the only – and not even the first – source of information."

- AI referred Sarah Mustillo, a former Dean of the Notre Dame College of Arts and Letters, to a promising sounding source entitled "AI and the Liberal Arts: Why Humanistic Perspectives Matter" (American Academy of Arts & Sciences 2023 report). Try as she might, though, she couldn't find it. When she asked AI why she couldn't find it, AI admitted "I **inferred a report title that sounded plausible but doesn't**

correspond to an actual, standalone published report under that exact name... I generalized the structure of [AAAS's] existing work into a citation that seemed like it should exist, even though it technically doesn't..." (Technically???) The former dean said her mind was blown, and she wondered if students knew that [AI] has a tendency to just make stuff up like this if they might be more careful about using it.

- An author asked AI to analyze her writings. I was extremely impressed by the depth and quality of AI's analysis. My mind quickly changed, though, when the author finally interjected "This is nuts!" AI's answers had absolutely nothing to do with her work. AI admitted that it hadn't actually read anything the author had written and was just guessing what the work was about! The author concluded "What ultimately transpired is the closest thing to a personal episode of Black Mirror I hope to experience in this lifetime."

I'll add that, even when AI does give correct information, it may just provide factual summaries and not be able to provide the analysis that is usually required in my course, e.g. it can't answer exam and discussion questions very well. I've asked AI questions I ask my students and students generally have better, more accurate, and insightful answers (at least when students have carefully gone over the material themselves).

III. So, is there no hope for AI???

No, that is NOT the conclusion reached by the MIT study cited earlier. In fact, it found that people who wrote without AI assistance first, then used ChatGPT to edit their work, actually showed increased brain connectivity. "That means: Do the cognitive heavy lifting yourself first, then use AI as a refinement tool rather than a replacement for thinking. Your brain needs that initial workout to stay strong when you later add AI assistance."

IV. AI can harm the environment

The MIT Technology Review warns that AI video may be especially problematic. "*When we tested the energy demands of various models, we found the energy required to produce even a low-quality, five-second video to be pretty shocking: It was 42,000 times more than the amount needed for a chatbot answer to a question about a recipe, and enough to power a microwave for over an hour.* If there's one type of AI whose energy appetite should worry you, it's this one."

V. Course Policy

If I could effectively ban all use of AI in my classes, I'd be very tempted to do so. But, I doubt that I could do that, so I am doing the following instead.

- There are several short assignments worth 1 to 3 points. Most of these involve taking notes on the readings and then discussing them in class. *Any and all use of AI on the short assignments is forbidden!!!* I want you to actually do the readings and take your own notes on them. If you make some mistakes or have things that you don't understand, that is fine; the small group and class discussions we have are designed to increase your understanding. If one student misses something, another may pick it up. I'm not super-

picky about the short assignments but whatever you do, I want it to be your work, not AI's.

- Incidentally, the readings have historically been one of the most popular aspects of the course. If you just have AI summarize the readings for you, you are going to miss out on a lot of very interesting material, and the entire class may suffer because of your being unable to meaningfully contribute.
- For papers and projects, you can use AI to identify sources that might be helpful. I do this myself. Further I suspect that knowing how to use AI as an aid to research is one of the things employers will be most interested in. However you must actually look at those sources and not just trust or copy what AI says about them.
- All first drafts of the major assignment must be written entirely by yourself. If you then want AI to refine your writing, you can. But, you must tell me that you did so, and include both your first and final drafts in your submission. In other words, exercise your mind first before you ask AI for assistance.
- Please, no AI-created videos!!! I don't want thousands of trees to die just so your paper or presentation looks a little cooler.
- You are responsible for the accuracy of everything you say. Sure, people make mistakes. But, *if your work includes obvious AI hallucinations, you will be severely penalized, and potentially even fail the assignment.* (In other words, if you try to make AI do most of the work for you, you better be very, very sure that AI got it right!)

Finally, I'll stress that this policy is for your benefit. You are here to learn, and not just get a good grade. Or at least that is my goal for you! But even if you are here just to get a good grade (perhaps with as little work as possible) your chances of getting that good grade are endangered if you don't learn how to use AI properly. Employers do not just want you to demonstrate that you know how to use AI. They want you to show that you know how to use AI properly, and that you've learned enough from your college courses to help you do that.