

ACADEMIC WRITING AND AI: A READER'S POLICY

Thomas Pepinsky
Cornell University
October 2026

Artificial Intelligence can generate good prose quickly and inexpensively. This creates opportunities and challenges for academic writers. AI will be a part of the academic workflow for the foreseeable future, but technological progress has moved more quickly than have social norms and expectations around AI use. The purpose of this note is to outline and explain my position on reading text generated via AI. This may inform your practices of writing text for any audience, regardless of whether you choose to generate it using AI.

This note is about academic writing. Academic writing includes dissertation chapters, article drafts, research proposals, referee reports, opinion pieces about academia, formal email correspondence, and related forms of writing. Most advice about AI and academic writing starts from the writer's perspective, asking how an academic should write and weighing the arguments for replacing some or all of that process with AI. These arguments focus on AI's effects on the writer. This focus misses the basic insight that academic writing is a form of communication, and the audience for academic writing is the academics who read it.

The biggest misconception for academic writers in the age of AI is that their readers do not care—or put more strongly, that they *should* not care—about how the text that they read has been created. This position is that text speaks for itself, and the audience's responsibility is solely to judge whether its contents are true or valid.¹ This is wrong, because academics in the human sciences almost never read for the sole purpose of determining whether text is technically correct.

¹ For a recent statement of this position, see Allison Stranger, "[Who Cares Whether Academics Use AI to Write?](#)" *The Chronicle of Higher Education*, October 13, 2026.

Instead, most of what academics in the human sciences do is read to provide feedback to authors. This is the main job of any teacher, adviser, mentor, editor, or colleague: to read drafts and reports and chapters and offer feedback on them. The purpose of reading these texts goes beyond the actual content of the text itself. Not surprisingly, academics who read every day become very good at reading, and develop strong views on what constitutes effective and appealing writing.

The feedback that academic readers provide focuses on the positioning, relevance, significance, clarity, style, and appeal of the prose as well as the validity and accuracy of its claims, because a teacher, adviser, mentor, editor, or colleague is reading for more than just the claims themselves. The process is subjective, and that is by design. As a teacher, my job is to help students to develop analytical and communication skills that are measured via their writing. As an adviser, I have a responsibility to support students who are independently capable of devising a research project, executing it, and communicating its results to a human audience. As an editor, I look for significant, creative, appealing, expressive, and clear arguments that are supported by coherent logic and judiciously interpreted evidence.

Reading for such purposes takes time. As a result, academics who read in any of these capacities face opportunity costs, because reading closely and offering feedback takes time that could be spent doing something else.

Writing is a medium through which humans communicate ideas, arguments, and evidence to one another. Using AI-generated text to communicate ideas and arguments to an academic audience is analogous to asking a fitness trainer for advice on how to lift weights when you are using a forklift. The process through which your sentences are constructed matters, because writing is thinking. Knowing this, you should reason through what a reader expects from your academic writing.

The first implication is that using AI to generate academic text breaks the direct link between the writer and the reader. In the language of pragmatics—the study of human communication in the social world—using AI this way is *uncooperative* if your reader expects to read human text.²

The second implication is that if you have sent an academic something to read in an academic capacity, you have asked a reader to give feedback to *you as an author* on the text that you have sent them. You have asked them to provide their time and effort, and you should respect that time and effort. Your reader has not agreed to read your academic text because they want to provide feedback about how Anthropic or Google is doing.

If you are a student submitting a draft, or an author submitting a manuscript, this is especially important. I do not mind reading imperfect writing by someone who is still learning how to write or perfecting their argument; that is my job as a reader. I am fully aware that AI can execute the task of writing and produce grammatically correct text. My time will not be spent evaluating AI writing.

A third implication is that you may find yourself producing text that does not have a human audience in the course of your work. Examples of this are computer code and scripts for statistical software. If you want to know if a function behaves properly or if a logical claim is formally valid, you should submit it directly to AI. Many services already exist that can execute such tasks, and some of them are free, and they are all faster and more accurate than submitting code and scripts to a human reader. If your academic writing is generated by AI, then the appropriate choice is to ask AI for feedback, not a human.³

² A linguist would say that AI text violates the COOPERATIVE PRINCIPLE, which says that you should communicate in the way that your interlocutors expect you to communicate. See H. Paul Grice. 1975. “Logic and Conversation,” in *Syntax and Semantics* (eds. Peter Cole and Jerry L. Morgan). New York: Academic Press, 41–58.

³ Gary King and Steven Worthington have recently argued that AI is demonstrably better at writing peer reviews than humans, by any conceivable metric. The reason that they do not conclude that journals should immediately replace peer review with AI review is because they understand the equilibrium effects of such a move. See Gary King and Steven Worthington. 2026. “[AI Article Reviews Dominate Human Peer Reviews at Their Best: Evidence and Policy Choices.](#)”

A fourth implication is that if someone has asked *you* to write something academic for *them*, you must tell them if you used AI to generate the text. There is no justification for hiding AI use in academic writing, except for if you intend to violate the cooperative principle, and you should not do that.

Fifth, that readers' expectations will vary for nonacademic audiences reading nonacademic texts. However, you should still remember the cooperative principle. If you want me to eat at your restaurant, your advertisements should show me pictures of your food, not AI renditions of it.

Finally, remember that writing has many purposes, but two are foundational in the academic context: writing is how you *think*,⁴ and writing is how you *communicate*. If you are sending your writing to someone like me, I assume that you are communicating to me about how and what you think. I will expect the text to be yours, and if I believe that it is not, I will not take it seriously. As someone who reads academic text every day, and has done so for decades, I have a good sense of what original academic writing looks like. I treat AI-generated academic writing as a waste of my time, and I will spend no time reading it.

⁴ Some people disagree, such as Joe Cruz, as he explains in "[I'm a College Professor. Writing Isn't as Important as We Think](#)," *New York Times*, Sept. 29, 2026. He and others are welcome to adopt different views about AI-generated text.