

High Standards for Critical Reading, Writing and Research in the Age of AI

Introduction

Recent advances in artificial intelligence (AI) have created a host of challenges and opportunities for higher education. The creation of the Penn AI Council, the AI Working Group in the School of Arts and Sciences, and a variety of smaller committees shows that Penn, like its peer institutions, is still developing an optimal approach to educating in the age of AI, one that promotes ideal faculty working and student learning conditions. We are the Critical Approaches to AI Working Group organized by the Price Lab for Digital Humanities, composed of Penn faculty, staff, and graduate students concerned about some of the impacts that AI is having on university education. As various units of the university consider AI policies for the classroom, the purpose of this paper is to make the case that Penn's approach should include innovations in AI-free instruction, specifically in areas where the new technical capabilities of AI models interfere with best practices in pedagogy. In particular, we focus on the humanities, and their investment in critical reading, writing, and research skills.

The Challenge

Though AI is a broad term encompassing many recent innovations, the primary challenge faced by humanities instructors arises from Large Language Models (LLMs) with generative text capabilities. Examples include ChatGPT, Gemini, Copilot, and Claude. These tools can summarize existing texts and generate new ones. Though they are prone to errors, like asserting false details (commonly called "hallucinating") or miscalculating math problems, they now come close enough to undergraduate work that instructors can no longer easily tell whether students have done work themselves, or whether they have used AI tools instead. This creates two important problems.

First, the widespread use of generative AI makes it impossible for instructors to evaluate student work and development.¹ There is no way to know how well students understood the reading if their summaries and syntheses come from ChatGPT, and no way to assess their development as writers and thinkers if their essays are generated by Gemini. Generative AI creates a barrier between instructors and students' thoughts, opinions, and logical and mental processes. It also can give students a false sense of their own skills and level of

¹ AI detectors are not a solution to this problem, because they do not work, they violate student privacy, and they are biased against historically marginalized groups, including non-native speakers. The use of detectors is therefore not a way forward; we must innovate our teaching strategies. See Liang, Weixin, Mert Yuksekogul, Yining Mao, Eric Wu, and James Zou. "GPT Detectors are Biased against Non-native English Writers," *Patterns* 4, no. 7 (2023).

understanding. Grading the work of the models is both a waste of faculty labor and useless for students, who have missed out on opportunities for both intellectual and social growth.

Second, students who use these tools may be developing skills of a kind—such as optimizing their prompting for current LLMs—but they are not developing as readers, writers, researchers, or communicators. In fact, they may even be getting worse. Members of the Critical Approaches to AI Working Group have observed a sharp decline in student skills, when we can be sure that they have completed assignments themselves. Colleagues at Penn and peer institutions report the same trend. Even some students are noticing a loss. At the Penn College of Arts and Sciences Fall 2024 AI Summit, an undergraduate panelist mentioned that many of his friends regretted that they could no longer write full paragraphs well. A recent magazine essay traces a Yale student’s perspective on her loss of learning and sense of authenticity as she started using LLMs to complete writing assignments.² Students in a writing seminar about AI writing, taught by a member of the Critical Approaches to AI Working Group, reflected at the end of the semester on the ways relying on AI for research, reading, and writing led to increases in linguistic homogeneity, as well as the outsourcing of their critical, creative, and emotional thinking.

Early studies support these students’ sense of skill atrophy. One study found that those using AI to write essays were overwhelmingly unable to provide any quotations from them when asked, and experienced cognitive decline relative to those writing without AI.³ In another study by a Penn faculty member, participants using AI to search for information developed “shallower knowledge” than those using older Internet search methods, with downstream negative effects on their ability to communicate in a convincing and creative way on the topic of research.⁴ A study by another Penn faculty member has shown that the unrestricted use of AI-tutors in high school mathematics harmed learning; when guardrails were placed on AI use, outcomes improved, but there was still no net gain over students who learned via traditional methods without AI.⁵

² Moore, Alex. “Inside Yale’s Quiet Reckoning with AI,” *The New Journal*, October 28, 2025.

³ Kosmyna, Nataliya, Eugene Hauptmann, Ye Tong Yuan, Jessica Situ, Xian-Hao Liao, Ashly Vivian Beresnitzky, Iris Braunstein, Pattie Maes. “Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task,” Pre-print, published by the MIT Media Lab.

⁴ Melumad, Shiri, Jin Ho Yun. “Experimental Evidence of the Effects of Large Language Models versus Web Search on Depth of Learning,” *PNAS Nexus*, Volume 4, Issue 10, Oct 2025, pg 316.

⁵ Bastani, Hamsa, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakçı, and Rei Mariman. “Generative AI without Guardrails Can Harm Learning: Evidence from High School Mathematics,” *Proceedings of the National Academy of Sciences* 122, no. 26 (2025).

Some of these harms arise because of cognitive offloading, the process by which outsourcing a task diminishes a person's ability to perform that task.⁶ Cognitive offloading is not harmful if the skill was unimportant to begin with: Many people used to know a dozen phone numbers by heart, but because this information is now stored in cell phones, those same people now know one or two, if any. This is a loss, but it's not very important.

Reading, writing, and research, however, are very important. When we talk about reading and writing, we are not referring to basic literacy, but to the critical methods at the foundation of all pedagogy, particularly in a liberal arts education, and even more particularly in humanities courses. Reading involves skills like identifying credible research materials, interpreting arguments, describing aesthetic and rhetorical features of texts, understanding how form contributes to meaning, and absorbing information. Writing involves skills like constructing an argument, crafting sentences and paragraphs, considering one's audience, developing style and voice, persuading, editing and revising, and conveying information effectively. In these senses of the terms, reading and writing have always been fundamental to higher education, and Penn's curricular requirements have always centered them.

As we face new challenges in teaching these skills, it is worth remembering why they matter so much: they are the cornerstone skills in developing expertise in all fields, not just in the humanities,⁷ and they are essential in and of themselves.

On the first point, critical reading and writing are essential for learning the actual content of a given field. Almost all of what we know in every field has been written down. Reading an essay or a book entails deep engagement with factual information, thoughtfully crafted rhetoric, and a long-standing scholarly conversation. Writing a response reinforces that knowledge and helps a student think through that scholarly conversation, perhaps even contributing to it or developing a viewpoint on it. Using AI to summarize the same work is, at best, akin to asking a roommate what it is about. Having ChatGPT write a response is akin to copying that roommate's work. This method might have some value, and might get a student through a class discussion, but it leads to precisely the kind of "shallow knowledge" that forecloses meaningful training and personal growth. Gaining expertise on virtually any subject requires reading well; making use of that expertise, or adding to it, requires communicating it to others.

⁶ Gerlich, Michael. "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," *Societies*. 15(1), 6, (2025).

⁷ Katz, Claire Elise. "Stolen Valor: How the Humanities "@ Work" Are Hidden in Plain Sight," *Public Humanities* 1 (2025).

On the second point, critical reading and writing are indispensable skills regardless of the specific fields in which they are applied. Although they are the special focus of humanities instruction, their use is clear everywhere. A doctor needs to be able to read the latest cancer research; if she outsources the reading to AI, the summary is likely to miss details like the uncommon side effects her patient begins to discuss in the clinic. A consultant needs to synthesize new information in a meeting and summarize it on the spot; he cannot rely on AI to listen and speak on his behalf in front of clients. An entrepreneur needs to be able to draft genuinely innovative business plans outside of another company's proprietary text generator. A lawyer needs to be able to craft a closing argument in response to what she hears in the courtroom. And in an age of widespread disinformation, much of it AI-generated, everyone needs to be able to evaluate writing for its plausibility and authenticity. If Penn students are to continue producing knowledge and new ideas, they need to strengthen precisely the skills that generative AI use weakens.

It might be tempting to think that the new capabilities of AI have diminished the importance of, or even replaced the need for, critical reading, writing and research skills, but, as these examples show, these skills are just as important now as they ever were. They are fundamental to human interaction and the functioning of society. The need for talented thinkers about interpretation, style, rhetoric, and summarizing will only grow as generative technologies proliferate. Indeed, AI companies have spent the past few years soliciting humanities graduate students and staff to produce sample text for them and to help evaluate the output of the latest models.⁸ To the scientists and businesspeople creating this technology, the value of the humanities has never been clearer.

The challenge, then, is not replacing those skills, but finding a way to continue to teach them. Calculators are a useful analogy. Students at the middle school level are generally expected to demonstrate mastery in arithmetic skills before they can begin using calculators as part of their day-to-day math classwork. Calculators have replaced mental math for many, perhaps most people, and surely most arithmetic done in the real world is offloaded to them (or their digital equivalents). Yet we still understand that arithmetic skills are important, that students should learn arithmetic habits of thought, and that they will be better mathematical thinkers—even better at using calculators—if they learn the basics before integrating the computational tool. The solution is as simple as requiring calculator-free instruction.

Generative AI is far more advanced than a calculator, of course, so it's no surprise that we encounter its pedagogical challenges at a far more advanced phase of students' careers. Yet

⁸ Members of the Critical Approaches to AI Working Group know graduate students who have been recruited for and participated in these activities. We withhold their names because the companies involved typically require onerous non-disclosure agreements.

the analogy holds. If we can teach students without AI, we can give them time to develop important skills, and instructors the ability to assess those skills. Later on, those students may go on to use AI in their everyday lives, as part of an advanced reading, writing, and research practice. They may even need instruction, including humanities instruction, in the best ways to take advantage of generative AI. But they will be better positioned to do just that, to thrive as skilled readers, writers, and researchers and to critically assess generative AI outputs, if they first develop and demonstrate mastery over reading, writing, and the synthesis of diverse sources of information. Given that generative AI is a more flexible, experimental, and powerful technology than a calculator, it is even more important to moderate its introduction in student learning, and to prepare students to use it responsibly and intelligently.

In short, advances in LLMs and other AI models threaten the existence of humanities instruction in general, and of reading, writing, and research skills in particular. Without the reading, writing, and research instruction that students have traditionally acquired in a liberal arts education, they will lack the basic skills necessary to make meaningful contributions in any field that involves knowledge or communication, within and outside of the humanities. If Penn wants to continue to produce knowledge and new ideas, we must find a way to teach these skills without AI interference.

Solutions

Although generative AI poses an existential threat to liberal arts education if left unchecked, we believe that it can be managed with decisive action. A successful approach would involve three components.

Requirement for some AI-free coursework and authentic assignments. First, Penn should confirm its commitment to high academic standards by requiring that students complete some coursework without the assistance of AI. As a matter of principle, students should not be able to complete a Penn education without demonstrating that they can do their own work. As a practical matter, as we have shown, some skills can only be learned and assessed without AI. This is certainly the case for critical reading, writing, and research, and likely applies to other areas as well, such as basic computer programming, math, and mastery of the content in specific domains. Whether they major in history, economics, chemistry, or something else, students ought to know the subject matter of their majors deeply rather than shallowly. With a requirement that students do some AI-free work, we can be confident that Penn can continue to produce alumni who rank among the most qualified in the world in every field.

This requirement should be embedded within the undergraduate curriculum, but it also requires that faculty design (or re-design) courses around authentic assignments that truly measure student engagement and skill (e.g., participation without the aid of devices, in-person exams, oral assessments, in-class writing, and in-person skill-building activities). Authentic assignments are especially essential for the writing seminar courses required of all undergraduates.

Transparency. Second, Penn should improve the clarity of its requirements surrounding AI. In part because this technology has developed so quickly, students and instructors have struggled to find common ground about when, where, and how AI tools can be used. Everyone would benefit from clearer, more consistent norms. Some forms of AI use ought to be banned entirely in the absence of clear and justifiable exceptions spelled out by instructors. For example, there is virtually no case in which a student should turn in an essay generated by an AI model under their own name as if it were their own work. More generally, course descriptions or syllabi should be required to include a simple indication of their requirements about AI use. For instance, “green light” courses could allow substantial AI use, “red light” courses could forbid it entirely, and “yellow light” courses could spell out specific *do*’s and *don’t*’s for students. Green and yellow courses should then provide guidelines for citing AI when it has been used. At present, students do not always know what they can and can’t do with AI in a given class, and instructors may not even know about all the ways students might want to use AI. Clearer guidelines and authentic assignments would remove the need for faculty policing, which harms teacher-student trust. They would also enable all parties to return their attention to education, rather than attempting to debate complex and rapidly evolving technical challenges throughout every term.

Humanities laboratories and AI-free spaces. Third, for all of this to work, Penn will need to invest in resources that enable AI-free education, especially humanities laboratories and AI-free spaces.

Humanities laboratories would be modeled after their counterparts in the sciences. Courses making use of these would entail a revised scheduling pattern with designated laboratory periods for completing reading, writing, and research in low-distraction, low-tech or tech-free, community-oriented spaces. The Hum Lab for the trial Kite Foundation courses run by the College in Fall 2025 provides an adaptable model for scaling up. The point of these humanities laboratories would be to provide support and structure for students to carry

out the existing work of the course in ways that facilitate their learning, rather than to add to the workload of the class.⁹

Where humanities labs would supplement specific courses, AI-free spaces would enable any student to work free from the distractions of AI tools, and any professor to assign dedicated reading, writing, and research time. They would be large, comfortable, communal spaces that might provide AI-free machines (for instance, laptops with reliable lock-down browser software), analog reading materials, and proctors who would circulate to make sure students were not distracted by tools on their phones or other devices. Such spaces might entail different affordances at different locations and times, including casual group discussion among students, concentrated reading or homework sessions, and open timeslots available for essay-writing, so that instructors could assign prompts to be completed in a proctored, AI-free environment. Creating these new places would require a collaborative discussion and brainstorming between different disciplines, including the design school and the humanities.¹⁰ They have the potential to become signature student hubs on campus, new and innovative centers at the heart of learning, where students can commune, concentrate, and think.

With resources like these, instructors would be better positioned to evaluate student learning, as well as their skills in reading, research, and writing. Both labs and AI-free spaces would support the use of “blue-book” exams that would enable instructors to assess whether students have acquired deep knowledge during the term—for instance, whether they have understood the assigned books and essays—and to evaluate student writing and research skills. Longer-form writing, meanwhile, involves advanced skills that cannot be captured in a single exam, from composing sections to crafting an argument to engaging thoroughly with source material. In labs, professors could work with students on writing assignments that go beyond the blue book.

These three threads intertwine. Instructors can write clearer syllabi if they can simply require students to write an essay in the AI-free space. Penn can craft clearer and more rigorous academic policies if they can simply require students to take so many “red light” courses. The AI-free spaces will become a vital part of the college experience and positively affect learning habits if students expect to spend time there as part of their ordinary coursework. And though violations would surely happen from time to time, strong

⁹ The History Department plans to pilot a weekly tech-free study hall in Spring 2026, staffed by graduate students. Seminar and course participants will be encouraged to attend these sessions to carry out research, reading, and writing tasks in a low-distraction, supportive communal environment. If successful, this pilot might provide a model for other departments and programs to follow.

¹⁰ Initial conversations between the Critical Approaches to AI Working Group and faculty at the Weitzman School of Design suggest significant excitement around opportunities to involve students and faculty in the design of these innovative AI-free spaces.

institutional norms would make them much rarer than they are now, just as Penn maintains healthy academic integrity despite occasional cheating.

By the same token, none of these policies will work without the others. Just imagine a student who wants to master writing skills in an English class under the current status quo. To the teacher, it may be clear that having ChatGPT write a paragraph is a violation of academic integrity, but the student may not realize that. The student may want to write everything himself, but feel that he risks his peers getting better grades if they all use AI. He may mean well but face overwhelming temptation as he works beside friends who are chatting with text generators to speed through their assignments. Meanwhile, the instructor may feel that rigorous AI-free requirements are, effectively, unenforceable, and that strict guidelines could harm enrollment. As things stand, it is up to students to police their own behavior and up to instructors to create their own policy, and neither happens very effectively. This is precisely where we need thoughtful leadership, clear policy, and a commitment to the demanding, high academic standards that characterize Penn education.

Conclusion

The pedagogical effects of AI extend well beyond the issues treated here. Universities must contend with new ethical challenges like carbon costs, student privacy concerns, intellectual property theft, and psychological harms. They must also try to take full advantage of exciting new opportunities in computer science, linguistics, philosophy, business, and even pedagogy. AI policy will certainly evolve in parallel in many directions.

Given the complex and fast-changing nature of the situation, there is much to be gained by acting quickly wherever the course of action is clear. As generative AI models become more ubiquitous, and if they continue to improve, the gap will grow between those universities that commit to humanities skills like critical reading, writing, and research, and those that allow them to wither. Employers and graduate schools will notice which institutions are graduating students who have genuine expertise, who can interpret complex arguments and craft their own on the fly, who can evaluate the source and motivation of an article, and on and on. Whatever small differences may be apparent between two users of ChatGPT pale in comparison to the difference between someone who offloads a skill and someone who masters it. Penn has a chance to lead in the development of pedagogy robust enough to survive and even thrive in the age of AI.

Signed by the following members of the Critical Approaches to AI Working Group: J.D. Porter, Emily Hammer, Whitney Trettien, Brent Cebul, Lauren Alcindor, Jajwalya Karajgikar, Kerry McAuliffe, Michelle Taransky, Stewart Varner