

Editorial: AI, Ghosts in the Machine, and AINT

Peter Shea, Editor-in-Chief, Online Learning

Professor, School of Education, University at Albany, State University of New York

As we bring you another issue of *Online Learning* (the third issue of our 30th year of publication, by the way), I wanted again to include some reflections on the topic of AI and academic integrity. As noted by many journal editors, we are seeing exponential growth in the number of AI-generated submissions to journals. Recently, an essay in the *Chronicle of Higher Education* by Mya Poe and Dylan B. Dryer (2026) noted that the influx of AI-generated papers was degrading scholarship, making it more difficult for authors of papers located in the “missing middle” to attract citations. The essay suggested that, in light of limitations on how AI functions, the literature reviews in AI-generated papers typically include references to a few older seminal papers and establish a superficial connection to this work. The AI-generated work then builds out the literature review with references to very recent papers published in the last two to three years, mostly in lower-quality open-access journals. This is due to the difficulties AI may encounter in gaining access to journals behind publisher paywalls. Because these AI papers are missing high-quality, authentic manuscripts published in the middle years, scholars who have done the hard work of conducting legitimate research and who may need to document their increasing impact through citations are struggling to do so. This missing-middle problem is also a telltale sign that AI was used in the generation of a paper, which should increase editorial scrutiny.

The woes of AI do not end there. Another recent story in the *Chronicle* outlined the next wave of AI-enhanced academic dishonesty. The author, Susan E. Ray, has spent years researching, writing, and conducting faculty workshops on this topic and believed she had found an AI-proof solution to academic integrity issues that also made student thinking and learning visible. Ray had students keep academic journals and versions of class assignments in Google Docs and could therefore monitor the number of edits and the increasing total time spent editing a document. In combination with tools that detected AI-generated text, she believed she had covered all the bases. However, a new entity appeared that promised both to “humanize” the AI-generated text and to take as much time as needed to edit it in Google Docs to make the versioning appear as though a human had typed the assignment. The AI arms race continues, and some faculty report that the only way to foil AI cheaters is to conduct in-class oral exams of students’ knowledge, a time-intensive approach that proves difficult to scale to large introductory sections. It might be possible to use collaborative learning and synchronous online environments to address this issue of scalability. One can imagine small groups of students on Zoom, tasked with collaboratively quizzing one another on their understanding of class content, perhaps using rubrics provided to all. This redistribution of assessment could make the idea of oral exams more scalable. The downside is that students generally despise peer evaluations.

This fall I will be teaching an asynchronous online class to pre- and in-service teachers in which we will be reading a new work by Cory Doctorow. In *The Reverse Centaur’s Guide to Life After AI*, Doctorow argues that, in automation theory, a longstanding metaphor exists: that of a

centaur. A centaur is a human who is assisted by a technology such that the human benefits in ways they might not in the absence of the technology. A centaur is a human using reading glasses, a hearing aid, or a riding mower, for example. A reverse centaur is a human exploited by technology. For example, Doctorow states that Amazon uses AI to monitor warehouse workers and delivery drivers during their work hours in ways that leave them worse off than if the technology were not used. One example is AI dashboard cameras that check to see if a worker's mouth is moving in violation of Amazon policies against singing on the job. Clearly, most people do not want to become the workhorse of AI. Relying on AI to make important decisions while allocating mundane tasks to fewer and fewer humans appears to be the hoped-for business model at work in the trillion-dollar investments made by AI model builders, according to Doctorow.

This raises thorny questions about our growing reliance on AI both inside and beyond the academy. Research conducted by Michael Gerlich (2025) raises concerns about the relationship between frequent AI use, cognitive offloading, and critical-thinking abilities. The results of the study with 666 participants indicated a significant negative correlation between frequent AI tool usage and critical-thinking abilities, mediated by offloading cognition (or productive struggle) to AI. Another smaller study conducted last year at MIT by Nataliya Kosmyna and her colleagues (2025) found that students using LLMs to assist with writing exhibited underperformance at neural, linguistic, and behavioral levels relative to students who used only their brains or just a search engine. The concerns about developing critical-thinking abilities, the centerpiece of higher education, seem particularly urgent.

The purveyors of LLMs are aware of these concerns and have been working on new tools to avoid the offloading of productive struggle in formal education settings. Grammarly, Google, and OpenAI are developing more Socratic versions of their AI tools, sold to schools as writing aids that will help students through assignments without doing the work for them. A writer for the *New York Times*, Dana Goldstein (2026), tried the products from Google and OpenAI. Initially, they seemed to work, but Goldstein quickly found that the tools were quite willing to provide polished “examples” of prose for a college essay assignment. Company representatives admitted that the tools were works in progress and would require more educator feedback.

We found the same problems last year when a doctoral student, Ed Beck, and I tried to support pre- and in-service teachers in using AI following recommendations from Ethan Mollick (2023) of the Wharton School. Mollick's guidance came in the form of prompts that teachers could use with AI in the role of mentor, coach, tutor, etc., and thus avoid the offloading of productive struggle. However, the teachers we worked with found that, like the efforts described by the *New York Times* author, the prompts resulted in interactions that could be gamed fairly easily. Asking the “tutor” to write an essay, for example, was a hard “no.” Again, asking for examples bypassed the support role and produced much the same result as asking for the essay itself. We are hoping to work around this problem with the help of experts working in the field who have addressed this issue with much more sophisticated prompts. One such expert is Anna Mills (2026), who uses Substack to share lessons learned from a grant-funded project, “AI Feedback Customized for Student Writers: The Updated PAIRR.” This is a tested library of prompts and instructions that can be shared and reused under a Creative Commons licensing agreement. The overarching goal is to enculturate educators to see AI as a legitimate partner, one in which students are

challenged and supported to work in their Zone of Proximal Development (Vygotsky, 1978), with AI serving as a knowledgeable peer.

Contrast this approach with current concerns (bordering on panic) about the use of agentic AI in higher education. See, for example, an interesting and provocative paper from Bozkurt and his colleagues (2026): “This paper explores the emergence of the ‘ghost student’: a digital surrogate created by the coupling of Large Language Models (the ‘mind’) and agentic AI browsers (the ‘body’). These entities are capable of navigating Learning Management Systems (LMS), engaging with content, and completing assessments with human-like mimicry, often rendering the actual learner’s presence optional” (Bozkurt, 2026, p. 1).

What are the conditions under which a student would elect to forgo all learning and add a “ghost student” to an actual course? Aside from the risks of detection and the clear violation of any principle of academic integrity imaginable, the student is choosing to offload all cognition (and indeed all responsibility) for the course to an AI agent. In this sense, the ghost student represents the logical endpoint of cognitive offloading: AI no longer assists the learner in completing the intellectual work but effectively removes the learner from it. It may represent something beyond even Doctorow’s reverse centaur: rather than subordinating the human to the technological system, the ghost student removes the human from the learning process altogether. These developments may appear to constitute an entirely new academic-integrity crisis, but I want to suggest that treating AI as “normal technology” offers a more productive way to understand—and respond to—the problem.

While I will leave a fuller examination of Bozkurt and his colleagues’ analysis for another time, the solution they seem to propose is a return to orality. While they acknowledge the challenges of oral assessments and the scaling issue, they do not acknowledge that an entire industry—cybersecurity—works on issues in this domain. While we must admit this is a difficult problem, other sectors (the financial industry comes to mind) have not given up hope. Although many organizations that control and oversee sensitive data constantly confront would-be hackers, the solution has not been a return to paper and pencils but rather the employment of tighter cybersecurity regimes. While it may appear hopeless to do battle with agentic browsers, the answer is not to cede the digital learning field to the academically dishonest (and I don’t think that is what Bozkurt et al. are advocating). We may need to investigate new forms of pedagogy, develop new policies, and invest in more sophisticated cybersecurity measures. While this may be a form of war, and we may wish to avoid it, decades of experience indicate that this is simply the new reality with which we must grapple. The need for a combination of new policy, pedagogy, and cybersecurity is more sharply defined by the rise of agentic browsers, but this is familiar ground for many outside of academia. This approach is grounded in the AINT perspective (AI as Normal Technology) of Arvind Narayanan and Sayash Kapoor (2025), outlined in both their book, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*, and their website, in which they write:

“Now that the collective initial shock of generative AI has worn off, there’s a need for structured ways to think about how AI’s impacts might unfold, instead of (over)reacting to each new technical capability or emergent social effect.”

Again, I don't want to leave the impression that Bozkurt and his colleagues are overreacting. Agentic AI does warrant our full attention, and their article does a commendable service in describing the enormity of the threat. I'm just not sure that the solution is a return to pre-AI days; rather, the authors seem to suggest a program of research on foiling ghost students through minimally technological environments.

However, the appropriate response to increasingly agentic AI is neither an unwinnable detection arms race nor the abandonment of digital assessment. Higher education instead needs an integrated integrity regime combining AI pedagogy that preserves productive struggle, new forms of assessment that verify learning, institutional and course-level policies that define legitimate AI partnership, and a cybersecurity research program aimed at verifying human presence. While AI is indeed a disruptive technology, the AINT framework acknowledges this and suggests that, as with previous disruptive technologies, we will need to adapt.

I want to thank the authors, Assistant Editors, Associate Editors, reviewers, copy editors, and the Online Learning Consortium for their hard work and support, which allow us to bring you the 23 articles in this issue. Also, a big shout-out to our Managing Editor, Carrie Miller, for always going the extra mile to bring you the important work in each issue of *Online Learning*.

References

- Bozkurt, A., Crompton, H., & Fell Kurban, C. (2026). The devil is in the det[ai]ls: AI agents, ghost students, and the crisis of verified presence in an agentic AI world. *Open Praxis*, 18(1), 1–12. <https://doi.org/10.55982/openpraxis.18.1.1145>
- Doctorow, C. (2026). *The reverse centaur's guide to life after AI: How to think about artificial intelligence—before it's too late*. MCD.
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), Article 6. <https://doi.org/10.3390/soc15010006>
- Goldstein, D. (2026, August 17). What do students lose when they stop writing? *The New York Times*. <https://www.nytimes.com/2026/08/17/us/student-writing-essays-ai.html>
- Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X.-H., Beresnitzky, A. V., Braunstein, I., & Maes, P. (2025). Your brain on ChatGPT: Accumulation of cognitive debt when using an AI assistant for essay writing task [Preprint]. *arXiv*. <https://doi.org/10.48550/arXiv.2506.08872>
- Mills, A. (2026, January 27). AI feedback customized for student writers: The updated PAIRR prompts. *Anna Mills' Substack*. <https://annamills.substack.com/p/ai-feedback-customized-for-student>
- Mollick, E. R., & Mollick, L. (2023). Assigning AI: Seven approaches for students, with prompts [Preprint]. *arXiv*. <https://doi.org/10.48550/arXiv.2306.10052>
- Narayanan, A., & Kapoor, S. (2025). *AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference* (Updated ed.). Princeton University Press.
- Narayanan, A., & Kapoor, S. (2025, September 9). A guide to understanding AI as normal technology. *AI as Normal Technology*. <https://www.normaltech.ai/p/a-guide-to-understanding-ai-as-normal>
- Poe, M., & Dryer, D. B. (2026, August 10). AI is ruining scholarship. *The Chronicle of Higher Education*. <https://www.chronicle.com/article/ai-is-ruining-scholarship>
- Ray, S. E. (2026, August 12). The next wave of AI-cheating technology. *The Chronicle of Higher Education*. <https://www.chronicle.com/article/the-next-wave-of-ai-cheating-technology>
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press.