2025, Volume 10, Issue 2

The teacher's paradox: Balancing learning and integrity in undergraduate biology

Keefe Reuther

University of California, San Diego

Intellectual rabbit holes are my happy place. In part, I became an educator and biologist so I could channel this tendency productively. When ChatGPT entered our cultural milieu, I was ecstatic, convinced that the barriers to curiosity were shrinking.

Which is also why a recent teaching experience was so demoralizing. I discovered approximately one in five students used AI to cheat on an assessment. I was grading a homework assignment designed to help first-year biology majors practice coding and data analysis. They had pasted AI-generated code but never ran it to verify if it worked. The evidence was right there: elegant, syntactically perfect code that threw immediate fatal errors when run. I suspected many more had used AI undetected. In that moment, I shamefully sympathized with every jaded lament about "kids these days."

Hundreds of hours spent designing scaffolded curricula suddenly seemed wasted. Even though the assignment was meant as a safe place to practice, students gravitated to the gilded AI fix: feed the assignment into AI, copy-paste code, and submit it as their own work.

As I sat there staring at untested, AI-generated code, a troubling question came into focus: By leveraging AI, am I empowering my students or making it easier for them to sabotage their own learning? This is not a novel problem. We've seen similar problems accompany the emergence of cell phones, the internet, and modern search engines. Every technological advance expands possibilities for both genuine learning and superficial shortcuts. The same tool that can deepen understanding for engaged learners can enable others to bypass the productive struggle that builds expertise. For teachers guiding diverse students—varying in motivation, prior knowledge, and external pressures—this means any powerful technology will inevitably be used in both constructive and counterproductive ways. That, in essence, is the teacher's paradox.

I believe that any classroom strategy for addressing this problem must include an explicit focus on developing fundamental science-process skills: asking clear questions, designing sound experiments, analyzing messy data, and utilizing technology wisely. This conviction led me to create, with my colleague Liam O. Mueller, a required course for biology majors on experimental design, data analysis, and computer programming. We piloted it just eleven months before ChatGPT changed

everything.

When generative AI went mainstream, I felt equal parts excitement and unease. It was another powerful tool that could either scaffold genuine learning or enable students to bypass it entirely. The question wasn't whether to integrate AI (my students would use it anyway) but how to align it with deeper learning goals.

However, my first ChatGPT exchange was less than impressive. When I asked it for recent papers on soil pH and microbial diversity, it confidently produced a plausible PLoS One citation that did not exist. When I pointed this out, it doubled down with more fabrications. I was frustrated, but I also saw an opportunity. This wasn't just a flawed tool, but potentially an ideal teaching instrument. I could use AI as a forum for students practicing to probe, verify, and remain skeptical of authoritative-sounding answers, the same habits needed to vet news, papers, or political claims.

The core value of AI lies not in its ability to info dump but in the low-stakes practice it affords for foundational competencies—precise question formulation, goal-directed problem solving, and rigorous evaluation of its outputs. The pedagogical challenge is to channel its imperfections into durable discernment.

Successes

At first, I simply had students paste prompts into public chatbots. Students shared and verified AI claims via rich discussions on course boards, normalizing skepticism. However, general chatbots like ChatGPT aim to please, not teach; students often copied full solutions and learned little. So, I built custom Socratic bots that asked questions, suggested approaches, and guided discovery. Beyond pedagogy, the custom chatbots addressed equity: all students received identical access to high-quality AI assistance, regardless of their ability to afford premium subscriptions. The transformation in my daily teaching was immediate. AI helped students interact with course materials adaptively, tailored to their specific questions. I spend much less time answering basic syntax questions, explaining fundamental concepts, or reiterating syllabus policies. Instead, AI enabled me to create scalable solutions that allowed students to self-solve their problems.

My students seemed to have a similar reaction. Over 12,000 total messages from 275 students with the chatbot generating approximately 2.25 million words of tailored feedback—all within one quarter. One student's reflection stayed with me: "I found the tool very valuable…it provided immediate feedback, without giving you the correct code or answer at first. It gave me the chance to learn and understand my mistakes before I even got to what I wanted to know."

This was exactly what I had hoped for. The AI wasn't doing the work for them—it was helping them do the work better. They were learning to ask better questions, to probe deeper, to think critically about sources and solutions.

Costs

It was amid that intellectual joy, when I discovered widespread cheating on multiple assessments. Beyond blindly copied code, student responses on asynchronous quizzes were surprisingly similar and oddly verbose—telltale signs of unauthorized AI use. Recent mixed-methods research shows some undergraduates may intend to cheat even when told it harms learning or could bring penalties (Huang et al., 2025).

My efforts to integrate AI were likely yielding both learning gains and losses (Wang & Fan, 2025). The painful realization wasn't the cheating; it was that I might be undercutting students' chance to struggle productively. Was I encouraging cognitive offloading rather than understanding?

My AI-inclusive curriculum helped intrinsically motivated students, but I was naive to think better tools alone would improve outcomes for everyone. Student decisions to cheat are shaped by multiple factors—including self-efficacy, achievement goals, perceived peer norms and behaviors, perceived likelihood of detection, and time pressure or other stressors (Bertram Gallant & Rettinger, 2025). I needed to redesign the course so the most productive, ethical path was also the most incentivized.

Design Response

My approach now centers on three integrated strategies that address the teacher's paradox directly. First, I have built a more intentional AI learning ecology—explicit activities, mini-lectures, and live walkthroughs that increase students' structured contact with AI, paired with crystal-clear policies about when and how to use these tools. Second, I have refined my custom Socratic chatbots to be more accessible, adding conversation starters and guided prompts for students uncomfortable with open-ended dialogue. The goal is to draw them toward pedagogically sound, course-specific tools rather than generic public chatbots. Third, I have restructured the relationship between formative and summative assessment.

My assessment strategy is based on the University of Sydney's two-lane framework. In lane 1, major summative assessments are secure and supervised—exams, interactive orals, or in-class demonstrations where AI use is controlled or disallowed. Lane 2 includes open, unsecured assessments where AI use is encouraged and scaffolded. Woven together, the lanes balance trustworthy evidence of individual competence with authentic opportunities to practice responsible AI use (Liu & Bridgeman, 2023).

Anchoring summative checkpoints in Lane 1 while moving practice into Lane 2, I implemented secure in-class midterm and final exams where the final grade could replace a lower midterm score—creating accountability while allowing for growth and mastery demonstration. Within the first quarter of using this strategy, the median score on the final exam was 15% higher than on the midterm.

Together, these three elements—intentional instruction, custom learning tools, and authentic assessment—create a framework where AI enhances rather than undermines genuine learning.

Looking back at this journey—from initial excitement about AI to the demoralizing discovery of widespread cheating to finding a workable solution—I realize the teacher's paradox transcends AI entirely. It's the eternal challenge of helping people develop genuine capability in a world full of shortcuts. Tools will evolve; the goal does not: helping students develop the intellectual courage to ask hard questions, the patience to struggle with difficult problems, and the discernment to distinguish durable understanding from ephemeral workarounds.

I remain excited about AI's potential to enhance learning. I still lose sleep pursuing intellectual rabbit holes, now often aided by increasingly sophisticated AI partners. But enthusiasm for this technology is meaningless unless we actively cultivate the foundational habits that make these tools useful. The best technology cannot replace our oldest one: a curious mind that asks good questions and refuses easy answers.

Institutional Review Board Approval

The included student quote was collected under IRB protocol #807175.

References

Bertram Gallant, T., & Rettinger, D. A. (2025). *The opposite of cheating: Teaching for integrity in the age of ai*. University of Oklahoma Press.

Huang, D., Hash, N., Cummings, J. J., & Prena, K. (2025). Academic cheating with Generative AI: Exploring a moral extension of the theory of planned behavior. *Computers and Education: Artificial Intelligence*, *8*, 100424. https://doi.org/10.1016/j.caeai.2025.100424

Liu, D., & Bridgeman, A. (2023, July 12). *What to do about assessments if we can't out-design or out-run AI? – teaching@sydney*. What to do about assessments if we can't out-design or out-run AI? https://educational-innovation.sydney.edu.au/teaching@sydney/what-to-do-about-assessments-if-we-cant-out-design-or-out-run-ai/

Wang, J., & Fan, W. (2025). The effect of CHATGPT on students' learning performance, learning perception, and higher-order thinking: Insights from a meta-analysis. *Humanities and Social Sciences Communications*, *12*(1). https://doi.org/10.1057/s41599-025-04787-y