

Chatbots Have Entered the Classroom. What Happens Next?

Columbia faculty are finding innovative ways to incorporate AI into coursework without compromising academic rigor.

By
David J. Craig
|
Fall 2026
Josie Norton

Getting your Trinity Audio player ready...

Last fall, Columbia students in an experimental new writing course fed short pieces of prose they had written into AI chatbots like ChatGPT, Claude, and Gemini and asked them to analyze and describe their own literary style. They then uploaded an excerpt of a stodgy academic tome — one that George Orwell had lampooned in his 1946 essay “Politics and the English Language” — and coaxed the system into altering the passage’s style and tone to resemble their own. They discovered that it often took dozens of carefully worded prompts to achieve just the right combination of, say, neo-noir and modern lyricism they were going for.

“The goal was to show that chatbots can be used to enhance writing, but that it’s not easy,” says [Dennis Yi Tenen](#), a professor of English who co-taught the course, Writing AI, with senior lecturer [Nicole B. Wallack](#). The course explores the possibilities and challenges of using chatbots to augment different types of writing. “You need to have strong fundamentals and a sense of your own voice, and then use

the tools surgically.”

Other students in a history seminar last year were tasked with using chatbots to analyze thousands of government documents, in search of hidden connections between US political figures and historical events. Their professor, [Matthew Connelly](#) '90CC, a pioneer of digital humanities who uses generative-AI tools in his own research, offered lessons along the way. “I tell students that I’m amazed both by what AI can do and by the outrageous errors it makes,” he says. “It can lead to discoveries, but also to bad scholarship, if you’re not careful.”

Computer-science students, meanwhile, were learning to use chatbots at nearly every stage of their design process, from writing and debugging code to stress-testing the software they’ve developed for all manner of industrial, biomedical, and cybersecurity applications. Along the way, they were graded on how well they used AI to improve upon their own manually produced code, as well as their ability to spot and correct mistakes in the AI-generated lines.

“Today, training students to be software engineers is like training symphony conductors, because on the job they’re going to be evaluating and coordinating lots of AI code,” says [Vishal Misra](#), a professor of computer science and vice dean of computing and AI at the engineering school. “We used to be training the violinists, the cellists. Now, they’re all conductors.”

Only recently have academics begun to talk so openly, even enthusiastically, about generative AI’s role in the classroom. Four years ago, when chatbots were still new and waves of US college students began turning in oddly competent yet soulless term papers and nearly flawless take-home exams, instructors at most universities (Columbia included) did everything in their power to stamp out the technology. They spent late nights uploading students’ papers into AI-detection programs and dreaming up other ways to spot abuse, such as planting nonsense terms in essay assignments and waiting to see if chatbots would faithfully regurgitate them in submitted works. Few professors even discussed chatbots in class, except to announce that anyone caught using one would be treated as a cheat.

Over time, though, the limitations of enforcement have become clear. As AI agents grew increasingly powerful, and as students became craftier at using them, professors struggled to detect who was relying on them and how. They found themselves spending inordinate amounts of time searching for traces of AI in

student work and feeling less confident in their conclusions. They worried about falsely accusing students of cheating and about creating an adversarial atmosphere in class, tarnishing everyone's experience. "No one gets into teaching to be a cop," says one English instructor who prefers to remain anonymous. "And if you focus too much on policing, it builds animosity and mistrust, which is counterproductive to everything else you're trying to achieve."

At the same time, business and industry leaders began to predict that workers in nearly all sectors, from finance and marketing to nursing and drug discovery, would in the coming years need to be highly adept at using generative AI. That's because the same large language models (LLMs) that power AI chatbots can, with skillful prompting, be made to perform a dizzying array of tasks, such as coordinating inventory audits and customer-service emails or predicting market swings and the disease-fighting potential of new molecules. "You're not going to lose your job to AI, but you're going to lose your job to someone who uses AI," Nvidia CEO and cofounder Jensen Huang quipped at a conference in Beverly Hills last year.

Meanwhile, scholars were beginning to employ chatbots in their own studies, even in fields that haven't traditionally been data driven, such as architecture, linguistics, and social work. To many students, it suddenly seemed apparent that knowing how to deploy AI in sophisticated ways was an essential career skill, whether they hoped to work in industry or in academia. Not surprisingly, many began to clamor for guidance in how to use the technology.

But if a central purpose of a college education is to teach students core knowledge and critical thinking, how could professors bring into the classroom a technology that might do that intellectual work for them? This is a question that academic administrators throughout higher education are now grappling with, and one that has no easy answers. Columbia's response has been to issue broad institutional guardrails around AI — the use of generative chatbots in classwork is acceptable only when explicitly allowed, for example — while encouraging professors to experiment. Faculty members say that the University, while strongly backing the use of generative AI in academic research, has given professors the latitude to establish their own AI course policies and test out new pedagogical strategies. Across the University, these experiments are now taking many forms — some aimed at limiting students' dependence on AI, and others at teaching them to use it more intelligently.

“We’re going through a period of remarkable change, and at times like this it’s important to keep an open mind about which strategies are going to work best,” says [Angela Olinto](#), the University’s provost. “We want to have lots of flowers blooming. One of our community’s greatest strengths is our adaptability and willingness to try new ways of teaching, learning, and conducting research.”

“What is this thing? What can we do with it? What does it mean for us, as researchers, as teachers?”

Before they could prepare students for success in the age of AI, most college professors had to get up to speed on the technology themselves, in what amounted to one of the most dramatic and rapid pivots in their own careers. Even at Columbia’s engineering school, where computer scientists like John R. Kender, [Kathleen McKeown](#), and [Salvatore Stolfo](#) laid the intellectual foundations for modern AI, and where the University’s influential [Data Science Institute](#) was established in 2012, many faculty members were shocked by ChatGPT when it arrived in late 2022. “I’d been working on LLMs for years and knew how powerful they’d become, and even I was taken aback by how easy ChatGPT was to use, and by the social implications,” says Misra. In the months following the tool’s launch, he took a lead role in educating Columbia’s senior leadership and faculty on its capabilities. “People who weren’t working in AI reacted pretty much the same way,” he says. “Everyone was like, ‘Whoa, what *is* this thing? What can we do with it? What does it mean for us, as researchers, as teachers?’”

[Amanda Irvin](#), executive director of the Columbia [Center for Teaching and Learning](#) (CTL), which provides training and support services to faculty, says that the collective mood she witnessed among instructors in the months following ChatGPT’s release was one of surprise and uncertainty. “People in academia had the same questions that everyone had,” she says. “Many wondered, What exactly is the value of human intelligence now? And what does AI mean for my own role, and for the role of a university?” Her staff began hosting large, town hall-style meetings where data scientists like McKeown and tech-savvy humanists like Wallack gave faculty tutorials on AI tools and how students might use them. “I think that everyone came out of those conversations feeling that higher education was actually more important now than ever, but that we had to be open to rethinking some aspects of how we teach.”

Even as many educators viewed AI chatbots primarily as a threat to academic integrity, some faculty at Columbia immediately saw potential to harness the tools

for their teaching. Foreign-language instructors like [Zhongqi Shi](#) '18TC, who teaches Chinese, were among the early adopters, recognizing that chatbots could be employed as virtual tutors, conversing with students and providing them detailed feedback on their grammar and pronunciation. Shi worked with technologists at CTL to design a chatbot specifically to help students prepare for job interviews with companies in China; it critiques their speech, and even evaluates their stabs at small talk for cultural appropriateness. Students “get a lot more feedback than I could ever provide,” he says. Similarly, chemical engineer [Chris Chen](#) '22BUS created a chatbot assistant that intentionally makes mistakes. In using it, his students learn to sniff out bad engineering assessments, “like they might need to do working alongside a distracted colleague,” he says.

Some faculty went further in their early AI experiments, designing lessons that aimed to mirror how students might one day use the tools in their careers. [Gregory Wawro](#), a professor of political science, began teaching his students to use chatbots to write computer code capable of analyzing the vast quantities of demographic and public-opinion data now available to scholars. [Kamel Jedidi](#), a professor of global business, encouraged his MBA students to use chatbots to improve the market analyses they conducted for real companies. And public-health expert [Gabriele Windgasse](#) '93PH trained students to use generative AI to quickly turn scientific reports about emerging health threats into accessible fact sheets. Along the way, these professors encouraged students to use AI judiciously, in ways that wouldn't interfere with their learning, and to be mindful of the long-term cognitive risks of delegating too much of their thinking to chatbots.

Many of these pedagogical experiments doubled as research projects, with faculty documenting how students responded to the new approaches. Education specialists at CTL helped the professors collect and crunch their data and then ran workshops promoting their efforts to colleagues across campus. In packed weekly sessions, CTL staff gave instructors ideas for how to reduce temptations for cheating (assigning fewer take-home tests; requiring multiple drafts on writing assignments) and tips for staying abreast of rapidly evolving AI tools: Students themselves are often privy to the latest products and have good ideas about incorporating them into lessons. And they worked on how to communicate AI course policies effectively to students, stating very clearly how and when chatbots could be used, if at all, and explaining the reasons.

Irvin says she expects that some of these endeavors will have an impact beyond Columbia. “Being at a major research institution, we’re well positioned to find scalable solutions to how and when to integrate AI into teaching, while keeping human intelligence and judgment at the center of everything we do.”

To date, faculty in data-driven disciplines such as medicine, engineering, finance, and environmental science have been the most receptive to students using generative AI in their coursework, because expectations of AI competency are highest in these fields. And graduate-level courses, with their emphasis on professional readiness, tend to place greater emphasis on AI skill development than undergraduate programs do. This fall, in fact, Columbia is launching a master’s of science degree in artificial intelligence (MSAI) to train people to design and implement AI systems in half a dozen fields, including the arts and architecture. The program, run jointly by Columbia Engineering and five other schools at the University, combines courses in data science with training in how to apply artificial intelligence in a domain of a student’s choosing — such as medical diagnostics, policy analysis, or the design of sustainable infrastructure.

“We’re trying to fill a gap in the workforce right now,” says Misra, explaining that there’s a shortage of people “who know how to build and adapt AI systems for specific industries in which they’re needed.”



Josie Norton

So far, every opportunity that Columbia has created for students to learn about AI has been wildly popular. In 2024, a small group of faculty led by computer scientist [Adam Cannon](#) launched AI in Context, an interdisciplinary course designed to give undergrads the kind of AI literacy that is increasingly seen as essential for nonspecialists. Rather than teaching students how to build AI systems, the course explores how people are using the technology and being affected by it. Students trace the field's history from its origins in the 1950s to the recent emergence of LLMs and explore some of AI's most pressing questions: why chatbots hallucinate and give biased results; whether AI companies should be forced to reimburse the writers, artists, and scholars whose content they trained on; and how AI may already be harming the environment and jobs. The success of the class, which fills immediately each semester, inspired the engineering school to launch an AI minor last year. "The new program, like the course, has a strong ethics component, which will help Columbia graduates make responsible decisions related to AI, no matter what field they go into," says Cannon.

When it comes to using AI chatbots in coursework, things get thornier. That's true especially at the undergraduate level, where students are still learning the domain knowledge and analytic skills they need to evaluate a system's output. The ability of chatbots to quickly generate answers to questions on almost any subject, whether ancient Greek philosophy or human anatomy, poses a constant temptation to students facing late-night deadlines or struggling through required courses. For this reason, professors still ban the use of generative AI on the majority of graded assignments. However, to help students become savvy users of AI, many faculty members have carved out portions of assignments where the use of chatbots is permitted — such as for brainstorming or research — or have designed stand-alone lessons that treat AI itself as an object of study. It's now common for professors in many fields to bring up a chatbot window on their lecture screen, type in a few prompts, and ask students to critique its mistakes or biases.

Efforts to keep students from over-relying on generative AI look different from discipline to discipline. In many science and math courses, where grades are based largely on in-class exams, professors are simply putting less weight on homework assignments, which students might use chatbots to complete. At the other extreme is computer science, where, in order to prepare students for coding jobs that already require intensive use of chatbots, Columbia faculty have increasingly integrated AI into their curriculum. And in the humanities, where writing has long been the heart of assessment, some instructors are redesigning their assignments so they're harder to automate, while talking with students about the intrinsic rewards of doing the work themselves. For guidance, many are turning to Columbia's Undergraduate Writing Program (UWP), whose affiliated faculty teach University Writing, a course taken by all first-year students.

"We provide professional development in what's known as writing-based teaching, which holds that writing isn't just a mode of assessment but one of the main ways we learn to think," says Wallack, director of the UWP. "In writing, students learn to tolerate the sensation of being uncertain, of trying hard to finish something but not being done yet, of being skeptical of themselves. If they stick with it, they learn to trust themselves as thinkers and discover their own style. There's no substitute for it."

Wallack and her colleagues are doing everything they can to make sure students don't squander the chance. For years, they've been honing a set of research-backed teaching methods that reduce rates of plagiarism and enhance learning overall.

These methods include encouraging students to develop arguments in their writing that have personal resonance; discussing their ideas with them one-on-one; and guiding them through multiple drafts and revisions. Now the UWP team is going further, finding innovative ways to demonstrate to students how the use of chatbots for writing assistance, even for seemingly small tasks like brainstorming or editing, can limit their creativity and cognitive growth. Some of the UWP's newest strategies are based on research done in the program. For example, writing instructors Kirkwood Adams '11SOA and Maria G. Baker have conducted studies on the rhetorical and stylistic limitations of chatbots.

"What our findings show, and what we try to communicate to students in compelling ways, is that chatbots are essentially 'status quo' machines — they reduce any idea or piece of writing down to its most generic, statistically average form," says Adams, who in 2025 coauthored a paper with Baker that offers one of the most detailed analyses of how chatbots write and respond to text. And while polished, anodyne prose may be perfectly suitable for corporate reports or emails, Adams says, it's anathema in the Core Curriculum, and especially in University Writing, a course that treats essay writing as a vehicle for developing the communication skills befitting scholars and thought leaders. "In this course, we're welcoming students into the academic community and preparing them to participate in it, to be fully conscientious and engaged, and to begin to contribute to the creation of new knowledge," he says, "not to regurgitate homogenized data."

To drive the point home to students, Adams and Baker have designed a series of classroom activities that dramatize how a chatbot's editorial influence can drain the life out of language. In one exercise, students are told to choose any subject they know intimately — the TV show *Stranger Things*, say, or the music of Taylor Swift — and to jot down a list of the most interesting, meaningful things they know about it. They then ask a chatbot to generate its own list on the topic. And they ask it to repeat the task, ten times at least, starting a fresh conversation before each prompt. In aggregate, the bot's responses are obviously generic and predictable, Adams says, and students feel a rush of appreciation for their own, more distinct insights. "They see that, if they ask a bot to write or edit something for them, they're allowing something special and valuable about themselves to be erased," he says.

So how are most Columbia undergraduates actually using AI these days? And how are they responding to their professors' efforts to adjust? Several studies are underway to find out. The most ambitious, led by Connelly, Wallack, and CTL

education specialist [Adam Brown](#), has for the past several semesters been anonymously surveying first-year students about how they use AI in writing assignments. As part of the project, the team is also conducting a randomized controlled trial to see if any of several new Columbia-developed teaching methods are reducing students' reliance on chatbots. The study is the largest of its kind in the country. The researchers, who have also surveyed Columbia faculty members about their perceptions of students' AI use, haven't yet released any findings, as the trial is ongoing. But they expect to do so in the coming months, and they say their preliminary data has already challenged some of their own assumptions.

"Our findings indicate that students are using chatbots in more varied ways than many of their professors realize," says Connelly, a history professor and the vice dean for AI in the Faculty of Arts and Sciences. "Very few students will use AI to generate a whole essay and just submit it. That's almost never the case." More commonly, he says, a student will recruit a chatbot to help generate ideas, locate sources to cite, or copyedit. Or in a pinch, they may use AI to summarize a book they haven't read, write an outline, or draft the conclusion to a paper. That's hardly ideal, Connelly acknowledges, but it paints a more complicated picture than many faculty have imagined.

Connelly says that students are also deeply conflicted about their own use of generative AI. "They recognize that it's a tool of great efficiency, and that they might need to learn it for work, but they're also concerned about how it could be affecting their learning," he says. "Many of them are actually choosing to use it less in their schoolwork, over time."

In other forums, Columbia students have said they wish they had more opportunities to speak candidly with faculty about the difficult choices they face, without fear of being suspected of wrongdoing. At a daylong symposium that Columbia organized earlier this year, called Reimagining Teaching and Learning in the Age of AI, student participants described some of the questions they and their classmates confront: If chatbots are banned in a course, does that really mean they can't use them at all, under any circumstance? What if they were to use them in sophisticated ways that enhance their learning, such as summarizing large numbers of articles before deciding which ones to read? Or engaging in a Socratic dialogue about the strengths and weaknesses of their work before submitting it? And what if their classmates are doing all these things and more, but they refrain — won't their own grades suffer in comparison?

Connelly says that whenever he has incorporated AI into classroom exercises, he's been struck by students' thoughtfulness in wrestling with the technology, as well as their pent-up eagerness to discuss a topic that, in some corners of academia, still carries a stigma. "Ever since ChatGPT was released, I've found that students are filled with legitimate questions about it," he says. "They're not all looking to take shortcuts. Many of them are just trying to make sense of the technology — to know when it's useful, when it isn't, and why. And they're looking to us, their professors, for guidance."

This article appears in the Fall 2026 print edition of Columbia Magazine with the title "Teaching the Chatbots."

Read more from **David J. Craig**



[Guide to school abbreviations](#)

[All categories](#) >

Read more from **David J. Craig**