Teaching Strategies in the Age of AI

In which some old ideas are polished and presented anew

(Revised 8.25.2025)

While there is clearly a role for AI in certain education contexts—particular assignments or even whole courses—there are also moments when we want students to roll up their sleeves and do the work we have designed for them. What follows is a list of practices that have proven useful in steering students away from AI when that is our goal.

A Few Preliminaries

Implicit in all of this is the assumption that students will make a good faith effort to do the work because they want to learn. To varying degrees, the practices below give them a foothold in their own thinking and writing (decreasing the likelihood of turning to AI in last-minute desperation) and/or make explicit the invitation to learn that is embedded in the work we assign. As much as possible, we need to avoid slipping into postures of constant suspicion. The practices below allow us to lead with our pedagogy and not with our skepticism.

A central part of the College's approach to AI is to let faculty define the norms and expectations for their classes. Students will thus encounter different AI policies in different classes. They may even encounter different postures toward AI in the same class. There is nothing new about different norms prevailing in different corners of the curriculum: it's no different than students learning that some classes require one kind of evidence to support a claim whereas other classes require a different kind, or that some classes (or assignments) allow for using the first-person singular in writing while others don't. As long as we are clear about the policies and expectations in our classes, students will be able to navigate the varied terrains of their education.

It's important to note that today's use of the term "AI" refers to widely used automation applications: decision making, classification, recommendation systems, and transcription/translation systems. "AI" is also used to refer to text and image generation tools. Strictly speaking, simply doing a search on Google or elsewhere is "using AI" because AI is increasingly embedded in those search engines, as well as in commonly used apps such as Grammarly and popular programs such as Word.

Our communications and conversations with students need to take these nuances into account. Monikers such as "Generative AI" and "Large-Language Models (LLMs)" refer specifically to the generation systems that have been fed large amounts of language (or code or images or what have you) in order to produce new texts (or code etc.) using prediction. Some refer to these systems as "synthetic media machines" or, very accurately, "synthetic text extruding machines."

Finally, note that many of these ideas (and some of the language) are lifted directly from documents circulated in previous semesters, including A Few Thoughts About Teaching & ChatGPT (rev. August 2024), Spring 2025 Resources on AI/ChatGPT (rev. January 2025), and Faculty Talking Points on AI (rev. January 2025). But this is new and improved. But it's old-fashioned, too.

A Collection of Useful Teaching Practices: What Faculty Can Do

1. Include an Explicit Policy About AI in the Syllabus: Students have made it clear that they want guidance from faculty. They are taking several classes, each of which will have its own stance toward AI. They may have friends who are using it with permission; they may have other friends who are using it illicitly and getting good grades as a result. They are besieged by ads and inducements to make their lives easier. We owe it to them to be very clear about the expectations and guidelines in each of our classes. A little class time and some clear language in the syllabus will go a long way toward framing this conversation. You can click here to see a small selection of recent AI syllabus statements from programs and colleagues at Bard.

2. Explain Why We Have Adopted Our AI Policy: As a rule, we all learn more from an activity if we know why we are doing it. This is one reason why highlighting our learning goals (sterile as that phrase might seem) is a sound pedagogical practice. A natural extension of such transparency is to explain why we have settled on our AI policy. In some cases, this will take the form of explaining to students what they will learn by engaging deeply with the class or the assignments we have designed. In others, it might involve explaining at which precise moments AI will be useful and why.

We should be unapologetic about the fact that, just as students learn different things in different classes, so too the AI policies in different classes will vary. Sometimes the work is difficult, and it is precisely by engaging with that difficulty that learning can happen. The following excerpt from an article by Corey Robin might serve as a useful prompt for writing and discussion about what's at stake in turning to AI:

> "There is nothing in the realm of work — no matter how interesting or exciting or desired — that does not entail, at some point, the experience of frustration, self-doubt, loneliness, and anxiety. Experiences that most of us (realistically, all of us) flee from, especially when we're by ourselves. Our goal shouldn't be to eliminate this discomfort. We need to teach students that it's part of the process, and develop strategies for coping with it. But for students to really get that — to believe it, to feel it — they have to do the work." (The End of the Take-Home Essay?" in The *Chronicle of Higher Education;* edited for brevity.)

Similarly, here is an excerpt from a recent NY *Times* article:

> "We tend to remember the satisfaction of having learned something, but we forget how difficult and challenging it can be to your identity to learn something new. To risk it, to fail, to take pride and want it to turn out great — all of those emotions are actually part of the whole process. And I will say, again, when A.I. takes out all of those emotional feedback loops that help learning happen it is actually not enabling learning. Just knowing a fact is not the same as learning, and that we are denying young people that experience of the pride of having acquired skill, talent and ability is, for me, just so sad." ("What AI Really Means for Learning" NY *Times*, 8.12.2025; if you don't have a subscription to the *Times*, you can get one for free via the library by clicking here.)

Either of these texts could be excerpted or assigned in their entirety if you want to devote more time to this topic. However you do it, a conversation in class along these lines promises to be more effective than outlining rules that we may have trouble enforcing. We should assume that students want to learn and then do our best to show them how the class is designed to engage them in that learning.

Many faculty have found it useful to clarify not just their policies but what AI is, what it does, and how it works. This can happen in the syllabus as well as in class, and it might need to happen more than once. Here are some previously circulated talking points along these lines.

3. Revisit the Design of our Assignments: We know from decades of pondering plagiarism that if an assignment is due all at once (e.g. end of the semester) and counts for an enormous percentage of the grade (i.e. high stakes for the student), the temptation to cut corners is huge. Faced with deadlines for several projects, some of which we have been procrastinating, few of us will make our best decisions. By contrast, a well-designed assignment sets students up to do their best work. Assignments that are scaffolded over time (broken down into parts or drafts) and graded with an eye on process (requiring and incentivizing the many steps along the way rather than valuing only the final product) have proven effective in inviting students to lean into their own work.

"Scaffolding" means different things at different times, but the underlying idea is to disincentivize the use of AI by breaking a larger, longer project into smaller steps (in class and out). Students are thus less likely to fall way behind and panic. To the extent that students resort to AI because they feel they cannot do what's being asked of them, a scaffolded assignment teaches them they can, in fact, do the work.

4. Revisit How We Ask Students to Read: AI's ability to summarize or offer opinions about texts is as tempting to students as it is daunting to those of us assigning reading to be done out of class. One way of adapting is to refine the ways in which we ask students to read. For example, teaching ways of annotating a text and assigning them to annotate a physical copy of the reading does two things at once: it teaches a way of engaging deeply with the reading, which is a valuable skill in its own right, and it results in a visible product – the annotated text – that can be collected. For a compact list of strategies to get students reading more attentively, see this list of 15 Reading Practices, which emerged from last year's Bard Reading Initiative Conference.

Could a student use AI for help in selecting which part of a text to annotate? Sure. But it adds a step to the process. And if we bring the readings into the next class and ask students to explain (in conversation or in writing) what they marked and why, we create a culture in which doing the work is rewarded. This points to what is ultimately the central point when it comes to reading: doing work in the next class session that students can only do if they have read with care.

5. Expand the Use of In-Class Writing: Asking students to write during a class session is a familiar practice at Bard, in part because it can achieve many different results. It can be used to refresh students' memories of what they read before class or what was discussed in a previous class; it can build on (and implicitly test) their understanding of readings or previously covered materials; it can deepen learning by asking students to apply material in some new context or in order to solve some previously unseen

problem. (If these ideas are new to you, I'm free for lunch most Tuesdays.) Beyond all this, it can help in the age of AI in a few particular ways:

- We get to hear students' actual writing voices, which gives us something to compare with the writing they complete outside of class. (This isn't fail-safe, since AI programs can mimic different styles, but it helps.)
- Prompts can be designed so that a student can only respond well if they have done the work assigned for a given day and/or if they truly understand the concepts being covered.
- In-class writing can be used to get students started on longer projects: to the extent that they have a foothold in their own thinking, students are more likely to continue down that road, since getting started is often the highest hurdle.

6. Use In-Class Assessments: Some faculty have taken the idea of expanding in-class work further by using in-class exams, oral exams, and/or hands-on projects that test students' mastery of course material. These are very much discipline-specific decisions, but what they share is the idea of rethinking how we use class time. The question we can ask ourselves (if we are in the habit of talking to ourselves) is, How could students show me in class that they have learned what I wanted them to learn? Not everything can be brought into the classroom, but much can. And while we tend to think of in-class assessment as primarily about grades, it can also (or entirely) be a means of learning what your students have learned: even if ungraded, in-class work provides feedback on how well they have mastered the material.

7. Embrace Experiential Learning: The invitation to rethink how we use class time is an exciting part of the current AI-infused moment, and these uses of assessment carry over to *any* kind of in-class work or hands-on project. The underlying principle is essentially that of a flipped classroom: using the time together to apply what has been learned (or read) outside of class. In any course we teach, it's worth looking for ways to get students working together to solve problems, to make something, to respond to something—in short, to *do* something—with the material they are encountering. To apply what they have learned, students need to actually understand the material. Knowing that there is such a task awaiting them will nudge any student who needs a nudge to take seriously the challenge of learning the material rather than pretending to learn it.

8. Use Group Projects: While not failsafe, having students work in groups tends to disincentivize the use of AI. As one colleague wrote last year, "students can sense very well when another student uses AI (perhaps even better than I can)." Group projects work best in conjunction with the kinds of experiential learning mentioned above.

9. Meet Individually with Students: Where class size and schedules allow, holding one-on-one meetings with students has many positive outcomes. In terms of AI, they offer us a chance to talk about work the student has submitted or work that is ongoing. We can ask students to explain their work and their process, and we can ask follow-up questions.

10. Make Use of In-Class Presentations: Much like a one-on-one meeting, presenting and explaining one's work in class requires that students understand (and can answer questions about) work done at home.

11. Require Portfolios: A portfolio requirement asks students to hand in their final drafts (of essays, lab reports, proofs, ethnographies, etc.) along with everything that they wrote along the way. These might include in-class free writes, photocopies of their notebooks, writing on napkins, as well as earlier drafts. These combine well with scaffolding and with process writing.

12. Add Process Writing: Whatever kind of work students are doing in our classes, we can ask that they explain their process of doing the work. Students will be familiar with process writing from L&T (and other classes), and it can take the form of a simple free write in class ("what's going well in the project that's due next week? What's proving difficult? What's your next step?") or a cover letter that narrates the steps they followed from the day they received the assignment through the day it was handed in (e.g. describe everything you did, right down to doodles on napkins, from the day you got the assignment to today.") By drawing attention to the process, we urge students to take it seriously and acknowledge that good work doesn't happen all at once. For more on elevating process over product, see this page from UC Berkeley.

13. Talk About Learning Goals: Perhaps the best way to deter students from using generative AI is to nurture the part of them that truly doesn't want to use it. Being clear and transparent, scaffolding assignments, and emphasizing process over product—all this will help. We can go further by engaging students in an ongoing conversation about what they are learning in our classes and how the assignments are supporting their learning. It has long been considered a best practice to invite students to identify (in week 1 or 2, say) what they hope to learn in a class. The presence of generative AI adds one more reason to do this: when students feel connected to what they want to learn in a class, they are more likely to hang in there during difficult assignments. Several of the ideas above have the effect of foregrounding learning and providing support for that learning, while at the same time acknowledging that the learning process is sometimes slow or difficult. This will always bring out the best in students as compared to an approach that emphasizes policing and enforcement.