

Artificial Education | The Harvard Crimson

Jill Lepore is the David Woods '41 Professor of American History and Professor of Law at Harvard.

This August, an MIT committee of faculty, students, and staff released a remarkably thoughtful report on “AI Use in Teaching, Learning, and Research Training.” Its many cautions included the warning that “as a community, what should worry us most is that many uses of AI deprive students of the opportunity to learn.”

By contrast, Harvard’s administration appears oblivious to such concerns — notoriously so. Harvard College Dean David J. Deming’s [recent recommendation](#) that student use of AI in writing-intensive classes should be not only accepted but encouraged, is only the most outrageous instance. This policy was, fittingly, the leading example cited in a Chronicle of Higher Education article titled “The Ivy League’s Surrender to AI Is Insane.”

Has Harvard indeed lost its mind? The comparison with the MIT report, set against a growing body of empirical research into the effects of generative AI on learning, is a good place to start answering that question.

While of course noting the astounding benefits of AI in many areas of academic research, the MIT report takes careful stock of the ways in which reliance on AI risks the abandonment of crucial educational experiences, including the problem set, the take-home exam, office hours, study groups, and undergraduate research opportunities. In the latter case, the committee found that instructors were considering outsourcing research assistance to AI instead of hiring undergraduates, presumably on the grounds of cost-savings and efficiency. The committee inquired, “But if those criteria come to dominate our decisions, we all have to ask, ‘What is it that we are here together to do?’”

Replacing research assistants with bots is a good illustration of the “[Waymo Effect](#),” defined by the coiner of the term, Daniel W. Hook, chief scientific officer at Holtzbrinck Publishing Group, as “what happens when a technology removes the friction of dealing with another human being, and we experience that removal as pure gain — because the costs of the friction were always visible to us, while its benefits were not.”

Getting a ride in a driverless car is convenient and fast, and if you don’t like talking to a driver or jostling with other people on the bus or the subway, it’ll seem preferable. But every time you take that ride, you’re missing out on encountering other people and, over time, you risk becoming more isolated, and lonelier, and less able to work with other people or, really, to even handle being with them. “Large language models,” Hook argues, “are the Waymo of intellectual life.”

At MIT, Harvard, or any other school, supervising an undergraduate researcher is a lot of work, costs a lot of money, involves a lot of meetings, and the results are not reliably excellent. Asking Claude is easier, faster, and may produce better results. It’s also cheaper, which is no accident: Harvard allows every Harvard student and faculty member to request a Claude token budget of \$2,500 per month, even as research funding is continuously threatened and the University [lays off](#) staff.

This approach, putting speed and efficiency by way of automation above all other considerations, comes not from the priorities of intellectual inquiry or scientific experimentation but from the business models of corporations. As Hook writes, “The problem is that we have made the human conversation the expensive option and the machine conversation free, and we are surprised at what researchers, responding rationally, then choose.”

But of course, the machine conversation is not free. It comes at the cost of our common life together as teachers, learners, and researchers, colleagues, mentors, and friends. Not to mention, supervising student

researchers is our job.

Considerably more is at stake than collaboration and community. Research into the effects of AI in education is preliminary but the results are, so far, grim.

A study conducted by Microsoft suggests that knowledge workers who rely on generative artificial intelligence engage in less critical thinking. A field study published by a lead researcher from Georgia Tech reports that “on average, participants in the ChatGPT group had worse learning outcomes than those using Google, especially for higher-order critical learning.” A study of high school math students led by researchers at the University of Pennsylvania found that students with access to ChatGPT performed better on tests, but when access was taken away they performed worse than students who had never had access.

Scholars have groped for what even to call this dire diminishment of intellectual capacity. “Human enfeeblement” has been proposed, and alike “cognitive surrender” and even “cognitive extinction.” (You could also call it, “losing your mind.”) And that’s not even accounting for the losses of trust and integrity. According to [a survey](#) last spring, more than a third of Harvard undergraduates polled used AI in ways that violated course policies. So far, and, to be sure, unintentionally, this College’s solution to that problem has been to relieve faculty of our obligation to teach and students of their obligation to learn and to make up for the attendant losses of discovery and humanity by allowing faculty to eat in the dining halls for free.

Deming said two weeks ago that he would produce advisory guidelines on AI use this year, taking faculty input into account. But nothing about how the administration has approached this issue suggests that this new policy will be meaningfully debated or even answerable to evidence.

What is it that we are all here together to do, again?

Jill Lepore is the David Woods '41 Professor of American History and Professor of Law at Harvard.