

Thoughts on AI and Teaching

By Jonathan Libgober

I have a memory from grad school of discussing how long problem sets were taking us during the first year of my PhD. I sensed I was spending less time on them than several of my classmates. For instance, I was able to take part in a number of activities which many others in my cohort did not have time for. But I was coming straight from undergrad with lots of math under my belt, so I was used to the grind of problem sets. What I remember most clearly is one classmate realizing that I was using Mathematica to do a lot of the algebra on our homework. He seemed disappointed when he found out that you could have a computer program simplify some messy calculations, work that had taken him quite a bit of time.

In May I listened to a commencement speech that argued AI would have a similar impact on work as the calculator (or, by extension, Mathematica). The analogy is comforting since it suggests AI is just another tool. But while I do think my experiences using tools like Mathematica influence my views about how students should use AI, I emphatically do *not* think that AI is just another tool. And more importantly, I cannot see how one could possibly know just how similar AI is to past tools students have used. Others have written about the [staggering degree of progress](#) in math and coding over the past few years. Well-informed people who have paid close attention to these developments have repeatedly [revised their expectations](#) of AI's trajectory. I cannot understand how anyone can predict how these tools will be used in the future with any degree of confidence in light of what we have seen.

The purpose of this note is to highlight how I think AI should be used in my courses and to provide a glimpse of how I am thinking about these issues. I do not think anyone right now has a complete, coherent philosophy of how to use AI in research and education, and that includes myself. I view teaching (and especially this semester's teaching) as a way to help develop my own philosophy of AI use. As such, my views about what constitutes good teaching and the right response to AI will be informed by how well students seem to grasp the course material. In terms of concrete plans, for now I am not going to change what I teach or the basic structure of class. But I am changing how I evaluate students slightly. And I am insisting that students consciously develop good practices. Taken together, these changes qualify as neither blanket encouragement nor discouragement of AI use.

The way I personally approach the topic of AI and education is closely tied to my perspective as an economist. I anticipate that many durable principles that have guided the training for several generations of economists will remain valid. I maintain this view even though I acknowledge that AI will require changes in how we educate students. This perspective is deeply connected to one of the eternal lessons of economics: to be cautious of "this time is different" logic. Just because each new development and technology has idiosyncratic features does not mean that familiar patterns fail to apply. The reason is not that our models are extraordinarily good (they are not), but because the focus on idiosyncrasies is ripe for overfitting and overextrapolation.

Simple, common-sense principles often remain relevant across otherwise very different environments, even when the full truth is not captured.

I do not think that AI will make the skills of economists obsolete. I also acknowledge that it's not too hard to imagine a world where that is in fact the case. Still, I attribute this ease to quirks of human imagination.¹ By contrast, I find that a lot of commentary on coming dramatic changes in society due to AI engages in a kind of magical thinking which ignores complexities that make predictions hard. My view is that we should not be afraid to acknowledge that the future trajectory of AI in academic research is unknowable *ex ante*.

My current PhD students might not see the irony in this view, but perhaps only because most if not all will lack vivid memories of the 2008 financial crisis. This event loomed large at the start of my PhD career. At that time, economists were blasted for failing to successfully predict such a massive disaster. Many viewed the failure to see these events coming as an unforgivable sin.

In the context of the financial crisis, economists may very well have erred. In fact, the profession may have done so in a way that made the situation worse relative to a world with greater amenability for stress-testing falsifiable predictions. Retrospective predictions always seemed to publish better, for reasons that would best be left to another essay. But my point is not to litigate the events of 2008. Rather, I am advocating for humble ignorance when it comes to AI predictions specifically, and mention the financial crisis only to highlight that this view should not be taken as vacuous.

In our current moment, refusal to assert a concrete prediction strikes me as the more virtuous approach. There are just too many known unknowns and unknown unknowns. And if there is no reliable precedent for understanding where we are going, and the outcome will depend on factors we cannot yet know, are we really better off making such predictions anyway? One power of economic reasoning is in identifying one's ignorance, and using that to at least provide a framework for resolving disagreements. For questions on the future of education and AI, such a framework is needed more than any shot-in-the-dark guess.

In my own work, such unknowns provide motivation for the search for choices that achieve favorable outcomes no matter what the truth is. My attitude toward AI in education borrows this perspective somewhat. With this as a backdrop, here are some principles in the context of AI and education that I expect to be helpful whatever future we end up facing:

- **Mastery is important:** To have a successful career, you will need to know why your research is valid, to the point of always being able to take ownership of it. That means being able to explain it, replicate it, defend it, and so on. This is true even though the value proposition for these activities might not be clear. This was always the case, but is coming into more focus now. Note that just because the value proposition is difficult to articulate, that does not mean it does not exist.

¹ After all, Grok has repeatedly assured me that this is unlikely to be the case.

- **Judgment and communication are as important as output:** Academia is about communication. A great idea that is communicated poorly is not one which will have a good shot of being influential. And your style of communication will be one of the most important parts of how you define your brand. Outsourcing this to any third-party will thus cost you something. Incidentally, at the moment, there are many aspects of writing which humans appear to do much better than AIs acting autonomously, although one could debate how true this observation will be in five years. Regardless, communication is so intrinsically connected to the academic project that it is hard to see how it can be separated from what we do.
- **Students will need to learn how to use tools responsibly:** Actively obscuring, lying or cheating to skirt the rules is per se bad, and is also bad strategy because it is never foolproof and the punishments if discovered are extremely severe. But failing to make yourselves familiar with the best tools available to do your best work is simply leaving benefits on the table. These benefits might be provided to you personally, but more importantly they could help society at large.

I know that students are typically more focused on finishing the semester at hand rather than acting in accordance with the arc of history or the evolution of academia. I used Mathematica as a first year because I wanted to make it through the courses as painlessly as possible, not because I felt that doing so would help secure my place in the canon of theory. But I do think it is important to impart just how much power students have right now. On a personal level, what happens this semester will greatly influence how I think about AI and education in the future. More broadly, we are at a time where old metrics for evaluation will be tested but new metrics have not yet been established. For that reason, I believe that the coming years may be one of the hardest times to be a graduate student in my lifetime.

This is all to say that you should remain vigilant, and treat this as the dramatic time that it is. In 2020, it was obvious to students and faculty conducting courses via Zoom that we were living through a unique time in history. While the signs are not nearly as overt, I think this cohort may eventually look back on the near future as even more distinctive and significant. Knowing very well that this could be wrong, I will design my class in a way that takes that possibility seriously.

I will not discourage the use of AI. In fact, I will actively encourage it, but only to the extent that it can be helpful and only in targeted, deliberate ways that supplement rather than replace old methods of learning. I continue to use many of the tools I used during the first year of my PhD for my research today, including Mathematica. AI tools are extraordinarily powerful, and some of you may find new ways to use them to accomplish things past (and perhaps even current) generations of economists could not. Those kinds of uses I believe should be encouraged. And one general principle I believe is that how you use research tools early on shapes how you use them for the rest of your career.

My own use of AI has ranged from inessential reassurance about wording, to help with calculations that transformed days-long tasks into hours, to enabling me to build tools I

otherwise never would have attempted. Recently I have come to think that in some cases I would have been better off not using AI and I have started to worry about developing overreliance myself. But on the whole I think the impact has been positive so far. On top of this, I am constantly trying to think of entirely new productive ways to use AI, although for the most part I have failed to come up with anything particularly non-obvious (so far). Most of the time I think my uses are standard. Still, I use it enough that it would be enormously hypocritical of me to tell students not to, and I simply do not believe prohibiting AI is the best way to train our students. In all honesty, I am not sure how I would have reacted to instructions to not use Mathematica on my problem sets. My feeling is that if we instead encourage students to honestly disclose how they are using AI, then they will use it more ethically and efficiently, and we will become better teachers. I will thus ask my students to do so whenever it seems relevant (erring on the side of excessive disclosure).

But even though AI use will be an important part of education, that does not mean I believe there are no dangers associated with this technology. Students should be aware that some dangers may be difficult to pinpoint yet become apparent later in their careers. I am probably worse at simplifying messy algebra than I would have been absent Mathematica. But I worry that the skills AI might inhibit when used in education are even more fundamentally tied to the value of the profession.

More concretely for today, I no longer trust problem sets as a proxy for hard-to-quantify components of understanding. I no longer trust that a student who never participates can still learn the material well enough for an exam by reading independently. The latter I have not figured out how to address; actively checking attendance still seems too “middle school.” But for the former, I will switch to in-class quizzes and graded presentations. To help prepare for these, AI might help with individualized explanations that are more effective than what I would be able to provide. For many problems, the ability to simplify algebra for free might have meant I gleaned more economic intuition from a problem than I would have without such tools. I can imagine that asking ChatGPT to clarify one impenetrable sentence in a textbook would similarly enable students to get the full picture when the alternative is waiting for office hours.

But AI in its current form still strikes me as very far from a substitute for the depth of understanding one is meant to receive from coursework at the university level. I do think asking questions in office hours builds some skills that simply using AI to resolve confusion does not. In terms of assessment, I cannot see how good performance on the new metrics could be consistent with worse understanding compared to students who completed problem sets in the pre-AI era. I encourage you to brainstorm ways to make sure that your understanding remains at least as deep as it would have been without AI assistance.

Part of me deeply regrets making such changes to my courses because there is a good chance I myself might have performed poorly on these metrics. Undoubtedly there are several students who would have performed well a few years ago and worse under the new system. But this is a different time, and nostalgia cannot be a reason to stay the course.

In the future, these changes may appear to have been either excessive or insufficient. But in my research that takes limited confidence seriously, the *typical* case is that the right action ex ante fails to be the right action ex post. That may well happen this semester. But the fact that another strategy might have been better in hindsight would not, by itself, mean that these changes were not the right ones.

My hope is that these changes will ultimately leave our department and the profession stronger and better. For instance, the changes we make in teaching may impart greater understanding to students than would have been possible before AI. I do think there are good reasons for optimism based on the scope of new abilities researchers will have at their disposal. But the real test will be whether this new cohort ends up learning the material better than previous cohorts. I think the answer to that question will matter a lot for where we are going, and I am looking forward to learning it this term.

—

Disclaimer: I wrote the first draft of this essay without any AI assistance. ChatGPT commented that my initial thoughts were too scattered and helped me produce a more linear version, which it then declared ready. Opus 5 then identified several remaining shortcomings, which I worked through with suggestions from both models. I continued to iterate, but at some point I gave up. My wife did one proofreading pass with several helpful comments, which I incorporated. I did not ask AI to check these.