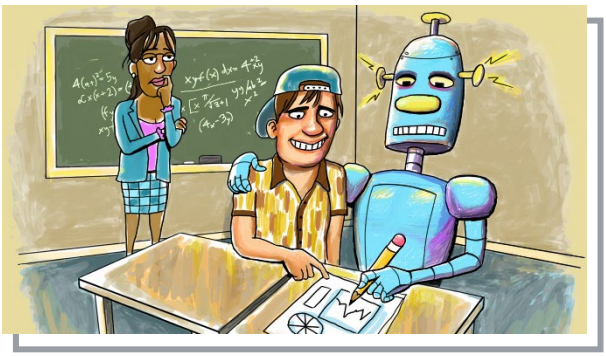
AI & Overreliance: Professor Peat's Dilemma

Professor Peat was an early adopter, and she jumped on AI like white on rice. A newly-minted assistant professor, Peat used AI tools to write her lectures and prepare her PowerPoint slides, to research and write her articles, and to review other professors' manuscript submissions to journals when she was called upon to do so. When others questioned her reliance on AI, Peat told them that she came from generations of farmers: "You know, once upon a time people had to dig up the fields by hand. Then someone invented hoes. Then someone invented horse-drawn plows. Then someone invented mechanical tractors. When new technology comes along, you've got to use it. Nothing wrong with that. It's better than living in the past."

But then Professor Peat ran across an article in the Chronicle of Higher Education written by Professor Chrysthanos Dellarocas, who teaches information systems in the Questrom School of Business at Boston University. He wrote, in part:

> "Last semester, a student stopped me mid-lecture. We were discussing how an online retailer should organize its customer and order data when he typed a few sentences into his laptop and turned the screen toward me. ChatGPT had produced a clean data model — customers, orders, products, payments — that looked remarkably like the diagrams I normally spend weeks teaching students to construct themselves.
>
> For a moment, my entire course felt strangely obsolete. But only for a moment. Then, I realized something even more troubling: Nobody in the room — neither the student who had generated it, nor any of his classmates — could say whether the AI-generated model was correct.

This article gave Professor Peat pause. She had just read an online news article about a lawyer at a prestigious law firm who had used AI to write a legal brief. Unfortunately, the AI – as such tools do



from time-to-time – hallucinated several nonexistent cases, which infuriated the judge and led to the attorney's client losing the case. The attorney was fired. Clearly, the lawyer's overreliance on the new technology proved costly to his client and to his career. This got Professor Peat thinking more broadly about the worry that some people have about "AI overreliance" and the negative impacts it might have now and on future generations.

Case study written by:

Robert Prentice, J.D.

Department of Business, Government and Society
McCombs School of Business
The University of Texas at Austin

Discussion Questions

[Note: the facts on the ground change quickly in the world of AI. You should do your own independent, updated research in answering these questions.]

1. Professor Carissa Veliz tells the story of two airplane crashes. One plane crashed because the pilots were chatting amiably as the plane landed, neglecting to monitor the airplane's speed and altitude. The other crash occurred because part of the airport's instrument landing system was out of service for repairs and the pilots could not remember how to land the plane without the use of high-level automation. "As flying has become more automated," she writes, "pilots have been losing certain skills required to fly manually." Veliz then argues that AI is the ultimate form of automation. Does this give you grounds for concern that people might over rely on AI, losing skills that are critical to their jobs, or perhaps just essential to human communication? Does your thinking change depending on the job or work area – for example, in aviation, in medicine, in legal professions? Explain.

2. Veliz gives another example: "When metals and electricity company Norsk Hydro received a ransom[ware] demand, they managed to avoid shutting down by going manual, but only thanks to older employees and other workers who returned from retirement to help out." Does this situation have relevance for AI and the industry that you intend to enter? Explain.

3. Using AI as she does, Professor Peat may never learn to write professionally and effectively on her own. Maybe she doesn't need to, since she has AI. But Professor Veliz writes:

> Among the abilities that writing improves is empathy. Writing asks you to put yourself in the reader's shoes. Will they understand what you are trying to communicate? Are you sure you won't bore them to death? Are you being respectful enough to other points of view? Social skills are some of the most valuable skills we possess; we would do well to make sure they are not lost to AI.
>
> If our students become used to relying on chatbots to write their essays, they might lose critical thinking skills, empathy, and creativity.

Do you agree with Professor Veliz that these are serious worries? Or are these concerns no big deal? Support your opinion with evidence.

4. In an article in The Atlantic, Rose Horowitch also addressed AI and its relation to writing skills:

> Cal Newport, a computer-science professor at Georgetown University, argues that the process of writing forces people to think in an orderly, linear fashion. It exposes flabby thoughts and shoddy reasoning. And the time and focus it takes to form thoughts into words, sentences, and paragraphs allow the author to make new connections and discover new insights.
>
> This feels true to me. My job is to write. … The prospect of a painful illness fills me with less dread than a blank page. But there's satisfaction in the struggle. The writing process is how I refine and formalize inchoate ideas and gain new understanding. By evaluating my arguments and discarding those that aren't convincing, I find the ones that are. Writing is hard because the writer is learning. If AI eliminates the challenge, it also eliminates the learning.

Do you agree with this sentiment, or do you agree with the title of Horowitch's article: "The End of Reading is Here"? Explain.

5. Horowitch supports her point of view by noting:

> Early studies have suggested that this is exactly what happens when people use AI to write. The process is easier. The product is often better than what someone could compose on their own. But it comes at the expense of mental development. One study in Brazil determined that undergraduates who used AI for studying performed significantly worse on a surprise test than those who studied without AI. The students trailed their peers even on questions that demanded reflection and effort instead of specific knowledge. Another

> study of hundreds of individuals in Britain found that frequent AI use for cognitive tasks is negatively associated with critical-thinking abilities.

Do these studies make you worry that AI "overreliance" may be a real thing? If not, why not?

6. Author Cory Doctorow (who wrote the imaginatively titled Enshittification), pointed out in a more recent book (also imaginatively titled--The Reverse Centaur's Guide to Life after AI):

> "A peer-reviewed study in The Lancet, a leading British medical journal, described a collapse in expertise by endoscopists who'd been assisted in interpreting colonoscopy results by an AI. These skilled professionals experienced a marked decline in their ability to identify potentially dangerous polyps after just months of working with an AI. This is classic "automation blindness," and bodes very poorly for the idea of using medical professionals as 'humans in the loop,' overseeing AI-based health diagnostic tools. "

Does the term "automation blindness" resonate with you? Why or why not? Can you think of a situation where you have suffered from it?

7. M.I.T. professors studied two groups of students who were assigned to write an essay. One group's members could use ChatGPT to help them, and the other group's members had to rely on their own brains. Brain scans showed greater electrical activity in the brains of the students who did not use ChatGPT. How do you feel about this? What does this signify to you, if anything? Explain.

8. The study in the previous question also showed that the vast majority of the students who used ChatGPT to write the essay could not, just a few minutes later, recall a single sentence that they had "written." Not so with the students who used their own brains to write their essay. Again, what does this mean to you? Does it surprise you? Why or why not?

9. Importantly, the M.I.T. scientists who conducted the study then had the students swap roles and write a paper on the same topic. The students who originally relied only on their brains to write the essay, but now could use ChatGPT, recorded the highest brain activity once they were allowed to use ChatGPT. However, the students who initially used ChatGPT but now had to write the essay using only their own brain power never did reach as high a level of brain activity as the other group. What lessons can we draw from this result?

10. A study by a Swiss professor found that higher usage of AI tools correlated with a decline in critical thinking skills due to "cognitive offloading," which happens when people reduce their own engagement in deep, reflective thinking by using external tools such as AI instead. Does this sound plausible to you? Why or why not? Does your own experience using AI provide support for or against the findings in the M.I.T. and Swiss studies? Explain and give your own examples.

11. Is it possible that AI could help your performance in one area of your work, while hurting your performance in another? A study by Fabrizio Dell'Acqua and colleagues involving numerous management consultants and three different tasks found:

> ...that AI capabilities cover an uneven set of knowledge work, and we introduce the concept of a "jagged technological frontier" to characterize the performance of AI when used by knowledge workers. Tasks that appear to human knowledge workers to be of similar difficulty may be performed either better or worse by humans using AI. Within this jagged frontier, AI can complement human work. However, outside the frontier AI is inaccurate, is less useful, and can degrade human performance. AI assistance improves human performance only for tasks within current AI capabilities—within the jagged technological frontier—and worsens human performance outside of it.

The key finding of this study is that when the consultants used AI in areas outside the model's competence, they performed worse than consultants who did not use AI at all on the same problems. The researchers characterized this as "falling asleep at the wheel." What do you gather from this study? Will you always be able to tell when a task is within the AI's frontier and when it is without? If so, how? Might you err by relying too much on an AI model that is ill-suited to the task?

12. In a very interesting study, Grace Liu and colleagues found that people working with AI systems quickly come to expect immediate answers and therefore to avoid working through challenges on their own. Therefore, as compared to people who did not use AI, these subject — although they do better in the short-run — will tend to perform significantly worse if they are denied AI assistance and tend to give up more quickly than those who did not use AI. They fail to develop persistence. Do the researchers' finding make sense to you? Why or why not? Is this finding consistent with your experience with AI? If so, how? If not, explain.

13. A study done by professors Abram Anders and Emily Dux Speltz found:

> "... that students may have spent less time putting ideas into sentences and paragraphs, but they put more effort into making higher-order decisions involving planning, judgment, evaluation, and quality control. One of the

students in the course had a great quote: "I had to learn how to think about my thinking." I wouldn't say that our study shows that more or less thinking was happening, but in a different way.

Students got really attached to their creative vision, and even when they were using AI, they really wouldn't settle for less than what they were going for, and it often led to a lot of back and forth."

These professors believe that AI can stimulate student creativity, especially if students critique the AI's product. How plausible does this sound to you? Is it consistent with your experience? Explain with examples.

14. This case study's mythical Professor Peat ran across an article in the Chronicle of Higher Education written by the real Professor Chrysanthos Dellarocas who noted that 1) reliance on AI can prevent students from learning important skills, and 2) the concomitant need to alter both how we teach and test students. Previously, his students were assigned a task (e.g., create a computer program) and the actions they needed to take to complete the task necessarily taught them what they needed to know. But now, his students can simply have AI create that computer program and can complete an assignment without learning anything. Professor Dellarocas explains:

    What do we want students to be able to do in a world of intelligent machines?

    The answer begins with distinguishing between two kinds of skill that AI has pulled apart. The first is artifact production: generating the essay, the program, the diagram, the analysis, the competency most of our learning outcomes have been built to develop, and the one generative AI has made widely accessible regardless of whether genuine understanding accompanies it. Bottom of Form

    The second is artifact reasoning: deciding what artifact should exist, directing a system toward producing it correctly, and judging whether the result answers the question it was meant to address. This is what AI cannot do. It is also what graduates are increasingly expected to do in the workplace: supervise, evaluate, and take responsibility for outputs that intelligent systems generate on their behalf.

    Artifact reasoning, in turn, has two components. The first is problem framing: specifying a question precisely enough that an artifact can answer it. Consider a question deans often ask: Which courses are most effective? Before any algorithm can answer it, someone must decide what "effective" means. Higher

exam scores? Greater improvement between the first and last assignment? Better performance in follow-on courses? Or something harder to measure, like the ability to apply ideas in unfamiliar situations? Each definition leads to a different dataset, a different model, and ultimately a different answer. An AI system can generate dashboards and rankings once the question is posed. But unless someone frames the problem carefully, the system will produce precise answers to a vaguely defined question.

The second component of artifact reasoning is critical interpretation: reading an artifact carefully enough to know whether it answers the intended question, what assumptions it encodes, and where it might mislead. Recognizing these failures requires genuine comprehension of the underlying system, the kind that only develops through sustained engagement with the material.

Taken together, these constitute the new core capability our courses need to develop, not as an add-on to existing learning outcomes but as the organizing framework around which courses should be redesigned.

My database course at Boston University's Questrom School of Business was where I worked this out. The course covers tools that are central to how organizations operate: data models, database queries, analytical dashboards. What changed was the question each course unit was designed to answer. In earlier versions: Can students produce these artifacts correctly? In the redesigned version: Can students reason about what these artifacts do?

How does Professor Dellarocas' solution sound to you? Can you create a similar solution for a course that you are taking or have taken? Or a job you have held? Support your reasoning with specific examples.

Sources

Abram Anders & Emily Dux Speltz, "Threshold Concepts for AI and Writing: Experimentation, Expertise & Agency, Computers and Composition, Vol. 81 (2026), at https://doi.org/10.1016/j.compcom.2026.10388.

Krzysztof Budzyn et al., "Endoscopist Deskilling Risk after Exposure to Artificial Intelligence in Colonoscopy: A Multicentre, Observational Study," Lancet Gastroenterol Hepatol, Vol. 10 (#10), pp. 896-903 (Oct. 2025).

Michael Coren "How to Stop ChatGPT from Ruining How You Think," Washington Post, July 7, 2026.

Alexandra Crosnoe, "Can Student Use of AI to Write Prompt Higher-Order Thinking?Chronicle of Higher Education, June 29, 2026.

Fabrizio Dell'Acqua, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality," Organization Science, Vol. 37, #2, pp. 403-424 (March-April 2026).

Chrysanthos Dellarocas, "When AI Can Do Everything, What Is Left to Learn?" Chronicle of Higher Education, April 13, 2026.

Cory Doctorow, The Reverse Centaur's Guide to Life after AI: How to Think about Artificial Intelligence before It's Too Late (Farrar, Straus & Giroux, 2026).

Len Gutkin, "Here Come the Student Resisters: The Tech Threatens to Make Young Americans 'into a Class of Drooling Morons," Chronicle of Higher Education, May 26, 2026.

Rose Horowitch, "The End of Reading Is Here," The Atlantic, Aug. 2026.

Grace Liu et al., "AI Assistance Reduces Persistence and Hurts Independent Performance," (2026), at arxiv.org/pdf/2604.04721v1.

Cyril Manning, "AI Gave Me Confidence I Hadn't Earned: The Dangers of Outsourcing Your Thinking to LLMs, The Chronicle of Higher Education, May 21, 2026.

Ethan Mollick, Co-Intelligence: Living and Working with AI (Portfolio/Penguin, 2024).

Helen Ouyang, "How A.I. Might Change the Way Doctors Think," New York Times, July 1, 2026.

Noam Scheiber, "We're Only Starting to Grasp the Pitfalls of Using A.I. at Work," New York Times, June 29, 2026.

Steven Shaw & Gideon Nave, "Thinking – Fast, Slow, and Artificial: How AI is Reshaping Human Reasoning and the Rise of Cognitive Surrender," (May 2026)(working paper).

Carissa Veliz, "Losing Skills," in AI Ethics (David Edmonds, ed. 2024).

Rebecca Winthorp, "What 370,000 College Essays Tell Us about A.I's Effects on Creativity," New York Times, May 27, 2026.

Tom Zeller, "Don't Quit This Whole-Brain Workout," Washington Post, June 23, 2026.