
The Return of First Principles: Education After AI

Odysseus Kaloidas
UC Berkeley

Abstract

Knowledge accumulates across generations, while an education must fit within a human life. Could AI help us recover facts and methods, acting as artificial memory while leaving more room to understand them? The opportunity is to spend more time on the concepts and assumptions that explain why a method works and when it applies. Easier access alone will not create that understanding. Outsourcing the reasoning can remove the practice needed to develop it, and unreliable answers still need checking. The question is which burdens AI can reduce while preserving the knowledge and habits that independent thinking requires.

1 An Education That Inherits Centuries

A modern student enters a conversation already centuries old. Before contributing, the student encounters its vocabulary, canonical results, methods, and exceptions. A bachelor's degree establishes breadth; a master's can add another layer of specialization. Some of this knowledge makes independent thought possible. Some effort goes toward retaining details that could remain accessible elsewhere. If AI made those details easier to recover, what should students spend the released time learning?

Jones's burden-of-knowledge analysis connects an expanding knowledge base with longer preparation, narrower specialization, and greater teamwork among innovators [7]. This evidence concerns innovators; it does not establish how much university teaching rewards memorization. It does, however, sharpen the educational question: how much of the inherited record must a person acquire before doing something new with it?

Antiquity offers an evocative comparison, with a smaller recorded scientific inheritance. It supplies no evidence that people generally memorized less or thought more. Plato's *Phaedrus* itself depicts the rehearsal of a speech and questions whether writing supplies understanding or merely reminders [10]. What makes the comparison appealing is the possibility of spending more time on fundamental questions while still drawing on the knowledge accumulated since antiquity.

AI could contribute to that possibility as artificial memory: reducing the effort of finding and recovering knowledge while leaving more room to examine what follows from it. Even a system capable of substantial reasoning can serve this role. A student may chiefly need access to what has already been learned.

2 A Library That Can Answer Back

People have long used records to extend what they can think about. Bush's proposed memex would let a person consult and connect stored material rapidly [3]. Clark and Chalmers argue that reliably available external information can become part of a cognitive process [5]. Risko and Gilbert review the related idea of cognitive offloading: changing a task by placing some of its demands outside the mind, as when writing something down instead of holding it in memory [11].

Generative AI extends this possibility through dialogue: a learner can request a forgotten definition, ask for the assumptions behind a result, or reconnect unfamiliar terminology with an understood concept. This shortens the journey from encountering a difficulty to finding relevant background. The learner need not know precisely which book or search term contains the answer.

As artificial memory, AI can still invent a reference, merge incompatible accounts, or retrieve the wrong passage. Checking and interpreting its answer consume attention too. The educational gain therefore depends on whether easier access releases effort for inquiry or merely replaces one demanding activity with another.

3 What Remains Worth Knowing

The prospect of retrieving facts does not make internal knowledge dispensable. *How People Learn II* explains how prior knowledge structures perception and learning [9]. Without relevant concepts, a learner may neither formulate a useful question nor recognize an answer's defect. Remembering and thinking are interdependent capacities.

That dependence gives first principles, the basic concepts and assumptions used to explain a problem, a particular role. Chi, Feltovich, and Glaser found that physics experts organized problems by underlying principles more often than novices, who relied on surface features [4]. A pulley and an incline can belong together because both can be understood through a conservation law, which identifies a quantity that remains unchanged. Knowing that relationship organizes many facts and identifies when a familiar formula is inappropriate.

Suppose AI supplies an equation of motion derived for a frictionless surface. The learner has saved the effort of finding it, but still has to decide whether it applies. If friction matters in the actual problem, the useful next step is to ask which force the equation omitted and how including it changes the motion. This is where access to a result can become an occasion for reasoning. Simply copying the equation skips that occasion, however quickly the answer arrives.

Conservation laws are one example of such organizing knowledge; basic definitions, causal assumptions, and standards of evidence are others. They do not allow every empirical fact to be deduced from a few axioms. Historical particulars and measured constants remain indispensable. Their educational value lies in helping a learner connect facts, recognize when a method applies, and understand why an answer might be wrong.

4 Available Answers, Acquired Understanding

A correct answer can come from understanding the material or from help that bypasses that understanding. Knowledge tracing distinguishes mastery from observed answers, allowing for errors despite mastery and correct responses without it [6].

Imagine that 60 percent of learners have mastered a skill and answer correctly 90 percent of the time. The remaining learners answer correctly 25 percent of the time. The overall correct-answer rate is 64 percent. If assistance raises only the second group's success rate to 50 percent, the overall rate rises to 74 percent even though mastery has not changed. These are illustrative numbers, not an estimate of AI's effects. The example explains why an improved assisted score leaves the learning question open.

The empirical record contains both possibilities. Bastani and colleagues found that general-purpose GPT assistance improved mathematics practice performance but harmed subsequent unaided performance; a learning-oriented tutor largely mitigated that harm [1]. Kestin and colleagues' randomized study of 194 undergraduate physics students found stronger immediate learning with a structured AI tutor than with the compared active-learning lessons [8]. Their different results make the design of the learning activity central to the question.

Westerbeek's August 2026 conceptual paper argues that cheap outsourcing can bypass the attention, retrieval, and judgment education is meant to develop [12]. This concern fits the present argument: time saved on accessing information and practice lost in understanding it have different consequences. Calling both processes offloading can obscure the difference.

5 More Knowledge Within Intellectual Reach

Bransford and Schwartz’s account of preparation for future learning recognizes that education can develop the capacity to learn from subsequently available resources [2]. That capacity becomes especially consequential when an extensive body of knowledge is available through conversation.

AI as artificial memory could make unfamiliar subjects easier to enter and connect. Whether that access turns into sustained inquiry depends on what the learner does with it. The possibility worth exploring is an education in which a larger intellectual inheritance leaves more room for asking why things work. Easier access makes that future conceivable; the habits developed through using it will help decide whether it happens.

References

- [1] Hamsa Bastani, Osbert Bastani, Alp Sungu, Haosen Ge, Özge Kabakcı, and Rei Mariman. Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26):e2422633122, 2025. doi: 10.1073/pnas.2422633122. URL <https://doi.org/10.1073/pnas.2422633122>.
- [2] John D. Bransford and Daniel L. Schwartz. Rethinking transfer: A simple proposal with multiple implications. *Review of Research in Education*, 24:61–100, 1999. doi: 10.3102/0091732X024001061.
- [3] Vannevar Bush. As we may think. *The Atlantic Monthly*, July 1945. URL <https://www.w3.org/History/1945/vbush/vbush.shtml>. Especially section 6.
- [4] Michelene T. H. Chi, Paul J. Feltovich, and Robert Glaser. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5(2):121–152, 1981. doi: 10.1207/s15516709cog0502_2.
- [5] Andy Clark and David Chalmers. The extended mind. *Analysis*, 58(1):7–19, 1998. doi: 10.1093/analys/58.1.7. URL <https://consc.net/papers/extended.html>.
- [6] Albert T. Corbett and John R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4:253–278, 1994. doi: 10.1007/BF01099821.
- [7] Benjamin F. Jones. The burden of knowledge and the “death of the renaissance man”: Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317, 2009. doi: 10.1111/j.1467-937X.2008.00531.x. URL <https://doi.org/10.1111/j.1467-937X.2008.00531.x>.
- [8] Greg Kestin, Kelly Miller, Anna Klales, Timothy Milbourne, and Gregorio Ponti. AI tutoring outperforms in-class active learning: an RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15:17458, 2025. doi: 10.1038/s41598-025-97652-6.
- [9] National Academies of Sciences, Engineering, and Medicine. *How People Learn II: Learners, Contexts, and Cultures*. The National Academies Press, Washington, DC, 2018. doi: 10.17226/24783. URL <https://doi.org/10.17226/24783>.
- [10] Plato. Phaedrus, n.d. URL <https://classics.mit.edu/Plato/phaedrus.html>. Benjamin Jowett translation, Internet Classics Archive; original fourth century BCE. Especially 228a–e and 274e–275b.
- [11] Evan F. Risko and Sam J. Gilbert. Cognitive Offloading. *Trends in Cognitive Sciences*, 20(9): 676–688, 2016. doi: 10.1016/j.tics.2016.07.002.
- [12] Hans Westerbeek. Designing for cognitive resilience: a distributed cognition approach to education in the age of generative AI. *AI & Society*, 2026. doi: 10.1007/s00146-026-03284-3. URL <https://link.springer.com/article/10.1007/s00146-026-03284-3>. Open Forum; published online 19 August 2026.