The Fundamentals Moved

The Classroom Should Too

Tyler Bell · September 27, 2026
catsplayingchess.com/writings/the-fundamentals-moved/

01

For years, I have been part of the same conversation about AI in education. It is a well-worn path that always hits common dead ends like "but students will cheat!" and "what about the fundamentals?" It has been nearly four years since ChatGPT was released, and the conversation has yet to meaningfully evolve. I understand why. Such objections are real and valid. However, these well-intentioned points often serve to preserve the status quo, subtly implying that the current system is working just fine. Just because our students earn degrees, does that mean they've learned all that we can help them learn? After they leave, are they "finally finished with school" or have we helped them become more curious, creative, thoughtful, and capable? I believe AI can help achieve this idealized goal of education.

02

At first, the tools were banned or heavily restricted. "Students need to learn the basics before they can use AI," most said. Over time, this position has become more difficult to maintain. AI use in courses has slowly become more prescribed, as instructors determine what aspects of their course students may or may not get help with. Students, of course, have been using it all along. And the use of AI is not likely to decline, especially as incoming university students will soon arrive with years of practice using AI throughout middle and high school (for better and worse). On the other end, we know that people will increasingly use AI systems in their careers. Therefore, if our goal is to educate and train technically excellent professionals, it is on us as their educators to help them learn to responsibly operate, and succeed, in this new environment. [1]

There is a natural tension here, especially when it comes to teaching, practice, and assessment. If AI helped produce the design, write the code, analyze the data, and draft the report, what does a submitted course artifact actually tell us about what a student has learned? [2] Can they explain the problem that was solved, the work that was done, and the assumptions behind it? Can they generalize what they learned by

modifying or improving the work if requirements suddenly change? Do the students understand their work, or do they merely possess something that looks complete?

To combat this tension, many have turned back to the traditional approach, such as paper exams, in-class quizzes, and handwritten essays. This faces the issue head-on by assessing student proficiency without access to AI, but it also (1) increases instructor overhead, potentially delaying feedback and limiting assessment touch points and (2) fails to incorporate the environment in which students will work, innovate, and live.

One response has been to adopt a form of oral examination. Sit down with the student. Ask them to explain what they've made or done. Can they vouch for the work, whether they created it independently or collaboratively with the AI? What happens if a constraint changes? How would their solution handle an edge case? Oral exams are an excellent opportunity to probe until it becomes clear what they understand and what opportunities for learning might still exist.

The oral exam seems to be one effective way to assess students while allowing AI use. We can personalize the examination and give specific feedback on a student's current strengths and weaknesses based on the work submitted and the subsequent conversation. Logistically, however, there are too many students and too few instructors and teaching assistants to do this optimally, especially in lower-level courses. In practice, each student may receive only two or three 10-15 minute oral exams per semester. As much as we might like to, we simply cannot talk with every student long enough to explore the complete depths of their understanding or the nuances of their misunderstanding.

AI has already started to change how students learn. It should also change what we teach, how we teach it, and how we think about the purpose of assessment.

03

Engineering has never required practitioners to perform every operation at the lowest available level. We teach hand calculations and first principles where they build intuition. Then we teach the use of calculators, solvers, coding, and computer-based design tools where the problem calls for them. Being able to perform every intermediate step is a skill, but it is not the ultimate measure of an engineer.

In my career, I've created software for many purposes and problems. But I have never had to go in and correct the assembly code generated by the compiler. Alas, I am a product of my generation: those before me certainly had to walk uphill both ways to a

program that ran well, writing assembly code by hand (or even punching their instructions into paper). By the time I was a student, the abstractions around solving problems with a computer had already earned a place in education and in ordinary work. The first programming language I learned was Java, then C++. Now, that does not mean I have no intuition for what happens below the language I am using. It simply means that I was able to learn and solve problems at a higher level first, only later studying what had been abstracted away.

To me, AI is the next layer of abstraction. Following this analogy, students should likewise have the opportunity to learn, think about, and solve problems at this higher level before studying what may now be abstracted away. In plain terms, that might mean students do not begin by writing code line by line in an IDE.

Anyone who posits that this may encourage an overreliance on AI is likely right. To that I would also say that we are then over-reliant on the compiler, the IDE, and the CPU. At what point do we earn the opportunity to take something for granted?

04

A few problems with this analogy arise immediately, of course. AI is a broader and less predictable abstraction than the tools that came before it. Modern AI can misunderstand requests, make assumptions, choose a fragile dependency, or deliver something convincing that works only under certain conditions. That does not eliminate or weaken the abstraction. It simply tells us where the work moves and what new expertise might be required.

In March 2024, I put it this way: “We will be able to write very sophisticated code very quickly with the assistance of these tools without ever thinking about or needing to go and look underneath the covers. It will just work. The large language model, in a way, is the new compiler, and I think that’s fine.” That said, verification becomes central. A student working this way may inspect generated code when the stakes or a failure demand it, but line-by-line review might no longer serve as the primary verification model. Consider what verification might mean at this level of abstraction, then.

- ☑ Do the tests pass?
- ? Wait, are the tests actually testing the right things?
- ☑ What could still be wrong even after all the tests pass?
- ☑ What is the scope of the solution?
- ☑ Under what conditions might it succeed or fail?
- ☑ Does the solution actually address the problem at hand?

Wait, what is the problem at hand?

As it turns out, the ability to answer these questions directly addresses the current shortcomings of this new abstraction. This is where the work and education should move.

Yesterday's fundamentals are no longer today's. And there is no reason to assume that today's fundamentals will remain tomorrow's. This does not mean the fundamentals disappear. It means that some of the practices we currently call fundamental are actually implementation details of the layers we currently teach.

This shift is already visible outside engineering. Abram Anders and Emily Dux Speltz analyzed final reflections from an experimental AI and Writing course involving 38 undergraduates across 22 majors. Students described learning to treat the writing process as iterative, bring their own expertise into dialogue with the model, and decide deliberately what to delegate. The authors frame the trade-off: "reduced cognitive effort in task execution demands increased metacognitive effort in planning and strategic oversight." [3] The work did not disappear; more of it moved into planning, evaluation, and judgment.

Beyond learning to use the tools, we need to teach people how to identify novel ideas and understand the complexity of the problems worth solving. [4] When AI can produce many plausible answers quickly, a distinguishing skill becomes curation. This involves exploring a broad range of ideas and knowing which ones to prune, not getting stuck on the technical hurdles of a single approach. The student who can generate ten designs but cannot explain which one should be built has not yet become a better engineer. I think that the best engineers will be the creatives, those who can identify and solve novel problems; the artists.

05

AI use should not be something students earn in their upper-level courses, only after they have learned "all the basics." Effective and appropriate AI use needs to be encouraged early on and even positively modeled by instructors. We naturally return to the earlier question: How do we ensure that a student is actually learning? The qualities that matter most at this new layer are also the ones a finished artifact reveals least clearly. You cannot always see discernment or judgment in a submitted artifact, especially one made with the help of AI. But you can hear it in a student's account of one.

Fortunately, AI changes both ends of education. In this AI-first approach, personalized instruction, practice, and assessment become scalable for the first time. The same AI systems that help a student learn and produce an artifact can ask the student to account for it.

> Q1
> Why did you choose this approach?

> Q2
> Which assumptions most affect your answer?

> Q3
> What happens if this part of the problem changes?

> Q4
> Which portions of your work did AI generate or influence, and how did you verify them?

An AI-mediated conversational examination (typed or spoken) can follow a student's answer, provide instant feedback at various levels of sophistication and granularity, and meet them right where they are. The system can map the evidence to an explicit rubric, probe the student to determine where their understanding thins, and return to a previous misconception the following week (or even the following day). Such a platform does not have to replace traditional assignments, lab work, design reviews, or human examination. AI-mediated learning simply helps make these forms of attention and feedback more frequent than staffing has traditionally allowed.

06

If each student has a different conversation, how can expectations and evaluation stay consistent? We do not need to standardize their conversations. We only need to standardize the learning objectives and the evidence required to demonstrate them (and many accredited degree programs already have similar structures in place). One student may need sustained questioning about their approach to a given problem. Another may need to defend a design decision or diagnose a failure. The questions and topics can adapt easily to each student while the required objectives remain consistent.

Traditional assessment practices will not suffice. The modern education system tells students that mistakes are part of learning and then averages those mistakes into their final grade. Industry and academia alike often treat that grade as a proxy for

understanding and capability. True learning and curiosity require somewhere to explore, have ideas, play, and be wrong. Importantly, then, the student-AI interaction history should not become a ledger of points for and against the student. Each course and task should have clear learning objectives and a rubric, and students should always be able to see their standing, inspect the evidence behind it, and know what they still need to demonstrate. If a student struggles early and later demonstrates mastery, the early failures should not remain deductions. This fluidity could also give motivated students more control over what they learn, helping them to push beyond the standards and assigned tasks.

And a course-level history is only the beginning of the possibilities. If such a record can support learning across one course, what would it mean for that history to extend across an entire curriculum? Or across all of education, work, and even day-to-day life?

None of this requires AI to become the final, unquestionable grader. The system's record of learning is a historian, not a judge. Consequential evaluations should remain inspectable and contestable, with human instructors and institutions retaining responsibility.

None of this makes cheating irrelevant, either. [5] Can a student game the oral exam by rehearsing it with another model? I hope so. A student who rehearses explanations of their own work, anticipates the questions they might be asked, and practices answering them is doing something we usually call studying. More importantly, in the ideal version of this loop, there is no discrete exam to game. Learning itself continuously produces evidence through explanations, revisions, and knowledge checks. [6]

07

Enough of the building blocks exist to begin developing and implementing this vision of education today. [7] A pilot could begin in a single course with defined learning objectives and a limited form of AI-mediated recitation. It could begin with formative questioning, progress to assisted oral defenses, and require instructors to compare the system's evidence and judgments against their own throughout. This could be done via a text-based system (chatbot-like interface) or real-time voice interaction. [8]

Lecture is a necessity of scale. One instructor addressing many students must speak toward some common denominator. Likewise, an exam must be practical to administer and grade. McLuhan saw the classroom as a factory almost sixty years ago,

concluding: “Mere instruction will not suffice.” [9] AI creates the possibility that lecture, practice, and assessment stop being separate phases and collapse into a single recitation-like system, all while encouraging creativity and curiosity. [10]

Done well, this could be the most hands-on education that has ever been available at scale. Students could receive more opportunities to explain and defend their work. Instructors could spend more time in the trenches with students, allowing their own expertise, experience, mentorship, and judgment to further shape the learning process. [11]

∴ **The fundamentals moved. We have enough to begin moving the classroom.**

Notes & Sources

1. “Environment” in McLuhan’s sense: a new medium does not remain a tool we pick up and put down; it becomes a surrounding condition that quietly rearranges habits, expectations, and institutions. AI is becoming that kind of environment for learning and work—which is part of why it has refused to stay politely inside course policies.

2. For introductory programming, the honest answer is very little. In 2026, Lepp and Kaimre evaluated five current AI systems against an introductory object-oriented programming course’s own assessments and grading criteria. The systems beat the average student cohort and frequently earned full marks on the longer tasks. Read the paper.

3. Anders and Dux Speltz offer this formulation while theorizing AI-assisted writing as cognitive work. Their qualitative study analyzed final reflections from 38 undergraduates across 22 majors in an experimental course taught in Fall 2023 and Fall 2024. It identifies patterns students described under deliberately scaffolded instruction; it does not show that AI automatically increases thought or validate AI-led assessment. Read the peer-reviewed paper or the Iowa State overview.

4. Terence Tao makes the parallel move for mathematics in “Mathematics in the Age of AI” (2026). Taking the arrival of these capabilities as a premise, he asks what the goals and values of the discipline then become. Read the essay or watch the ICM lecture via the Simons Foundation.

5. Strömberg, Lei, and Wu tracked secondary students after AI adoption. Homework scores rose roughly 18% while closed-book exam performance fell roughly 20%. The losses concentrated in students whose usage patterns looked like outsourcing (unusually fast, unusually high homework scores). Students who kept normal time on task lost little. Read the working paper.

6. Practically, there is a fine line between continuous interaction logging and continuous surveillance. Students must (1) know when they are practicing and when they are being evaluated, (2) have access to their records, (3) retain the ability to correct or challenge them, and (4) receive meaningful human review. Models must be evaluated for inconsistencies and bias across disciplines, communication styles, accents, languages, and disabilities.

7. The ideas I outline here are just that. They are neither perfect nor all-encompassing. Education is a human pursuit, and AI will not solve all of our problems. Human instructors, traditional practices, and tools that do not yet exist will all shape whatever direction education actually takes. We cannot stay still or wait for perfection. The tagline of the Science Center of Iowa's AI initiative resonates strongly with me: "Stay Curious—Stay Human."

8. Emerging modalities such as real-time video models, personalized AI personas, and world model simulations will also make learning more engaging, immersive, and accessible to a wider range of domains.

9. Marshall McLuhan and Quentin Fiore, *The Medium Is the Massage* (1967).

10. This could also reduce the extent to which two midterm assessments ultimately determine a student's entire grade and, by extension, their future opportunities.

11. Whether that surplus of attention is spent on students or spent on larger sections is an institutional choice; the technology does not make that decision.