

The Illusion of Learning

AI can make a student's work better before it makes the student better.

ENTRY № 07 · 22 September 2026 · 9 min read · From the series *The Next Era of Education*

Schools, in the form most of us went through, are young. Not the idea of teaching children, which is ancient, but the broad architecture we recognise today: a building, a bell, thirty children of the same age in one room with one adult, a curriculum that says what they should know by the end of the year, and an exam that decides what happens next. That architecture is only a couple of hundred years old.

Thinking about the future of education forced me to look harder at the past. If AI can respond to learners individually, it raises a more basic question: why is the system it is entering so standardised in the first place? I thought I knew the answer.


Industrialisation needed a workforce. Factories needed people who could read instructions, work to a clock and fill specialised roles, so we built institutions that could produce those capabilities at scale. It is a neat story. It fits the architecture of almost every school I have walked into, and it is the one many of us carry around.

It is also incomplete. When I went back to check it, I found that the part it leaves out is more interesting than the part it gets right. More importantly, it leads directly to the problem AI is now beginning to expose.

This is the first of a series. The frame for it is [The Five Eras of Education](#), which the diagram below summarises. This essay stays mostly in the third era, AI-assisted education, where responsive help is being added to the school system we already have.

Five eras. One expanding ambition.

Each era opens new possibilities. None automatically resolves what the previous eras left unfinished



Broad historical developments
 Where the AI layer sits today
 Possible directions, not destinations

Each constraint we remove reveals the next one.

Schematic: expanding possibilities, not a scale of educational achievement.

What mass schooling solved, and what it cost


Schooling predates the factory by a long way. What is new is mass schooling: the state-organised, compulsory, everyone-in-the-country version. When the comparative sociologists went looking for what drove its spread, the factory was not the best explanation. Their work on [the origins and expansion of mass education](#) and on [its spread across the world](#) found that it travelled with the modern nation state. Countries built schools for religious reasons, for political reasons, to make citizens who spoke the same language and told the same story about themselves, and yes, to make workers too. Countries with very little industry built them just as eagerly as countries with a lot.

What the factory story gets right is not the motive but the constraint. Whatever the reason a country wanted to educate all of its children, it faced the same problem: far more children than expert adults. That is a coordination problem, and much of what I

recognise as “school” can be understood as a solution to it.

The architecture of mass schooling

Mass education solved a coordination problem by making age decide everything after it



A powerful coordination system. But one in which time is largely fixed while learning varies.

Age sets the grade. The grade sets the curriculum. The curriculum sets the class, the class sets the timetable, the timetable sets the exam and the exam sets what happens next. Each link is sensible on its own. Together they let one adult teach thirty children, and they let a qualification mean the same thing in one town as in another. That is a real achievement, not a by-product.

The cost is individual attention. Picture thirty learners on fractions. Some get it in the first ten minutes. Others need one more example. One has missed an earlier concept and is quietly lost, and will stay lost, because tomorrow’s lesson assumes today’s. Teachers notice all of this constantly. Responding to all of it, every lesson, for thirty people, is beyond what one person can sustain. So the timetable decides, and the timetable is the same for everyone. Time is fixed. Learning varies.

I want to be careful not to make standardisation the villain. Common instruction is not merely administrative convenience. Well-designed collective teaching, [active learning](#) and discussion can be extraordinarily effective. The goal is not to dismantle the thing that educated the world. It is to keep its reach and its shared foundations while shrinking the gap between what a child needs and what the timetable can give.

Explanations everywhere

The first thing to loosen that constraint was not AI. It was the internet, and the clearest example is [Khan Academy](#). Long before generative AI, it had a serious go at the fractions problem: explanations you could watch at your own pace, pause, and replay at eleven at night before the test, exercises that told you whether you had understood, a dashboard that tracked where you were. The idea was older still. Benjamin Bloom’s [1968 account of mastery learning](#) argued that goals could be held constant while time and support varied between learners. What was missing was not the idea. It was a way to act on it for millions of learners at once.

Digital education made explanations abundant. What it did not solve was the harder problem: knowing what this particular learner misunderstood, and what should happen next. A learner who watches the fractions video three times may still not know which of the four possible misunderstandings is theirs.

AI may.

Help that responds

This is the moment we are in. AI is being layered on top of the existing system. The school year, the curriculum and the exams are all still there, and the new tools help students and teachers cope with them. [Khanmigo](#) is the obvious example, designed to guide learners through questions rather than hand over answers.

What is new is that the help responds. A well-designed system can ask a diagnostic question, offer a different explanation, pick a relevant exercise and check the answer. Take three learners who make the same algebra mistake. One is shaky on fractions. Another misread the question. A third just slipped on the arithmetic. A useful tutor would not assume the same wrong answer needs the same fix. It would form a hypothesis about which learner it is talking to, test it, and change its mind when the evidence changes. That is responsive teaching. Self-paced access, which digital education already gave us, is an important but more limited form of adaptation.

For the first time, then, responsive attention to the individual may become economically available at mass scale. Wealthy families have always been able to buy it, and good teachers have always differentiated where they could. AI does not invent personalisation. It changes its economics. And that creates a new problem, which is the subject of the rest of this essay. The system can now help a learner produce better work without necessarily helping the learner become more capable.

The illusion of learning

An assistant can improve the work a learner hands in without improving what the learner understands. A polished essay or a finished worksheet can hide the absence of independent reasoning, and it hides it from everyone: from the teacher who marks it, from the parent who sees the grade, and often from the learner, who has just had the pleasant experience of producing good work and reasonably concludes that something was learned.

This is not a hypothetical. In a [randomised experiment with nearly a thousand high-school mathematics students](#) in Turkey, a general-purpose AI interface improved practice performance markedly. Then the tool was taken away and the students sat an

unaided exam. The group that had used it scored about 17% lower than the group that had not. A tutoring version built with safeguards, giving hints rather than answers, largely prevented the damage. It did not produce a gain on the unaided exam either. That finding concerns those tools under those conditions, not all AI assistance. But it shows, clearly, that better output and better learning can move in opposite directions.

The reason is structural, and it reaches well beyond schools. When I am working, I want AI to remove effort. That is what a productivity tool is for. When I am learning, removing effort may remove the very process that creates the capability. The effort was not an inefficiency in the process. It was the process.

A LEDGER PRINCIPLE

A productivity tool removes effort. A learning tool must know which effort to preserve.

That is the first great design challenge of the AI-assisted era, and it is a different challenge from getting the answers right. A learning tool has to behave differently from a helpful assistant. It asks for an attempt before it helps. It gives graduated hints rather than a solution. It asks you to explain your reasoning, and sometimes it withdraws altogether. Assessment has to look at what the learner can do without the tool as well as with it, because both matter and neither substitutes for the other. The aim is not to preserve every old exercise unchanged. It is to know, for each task, what capability it is meant to build, and to protect that.

The same applies to teachers' tools. A dashboard that says a pupil is struggling with fractions is a prompt to go and look, not a diagnosis, and the teacher still has to decide whether the suggested fix makes sense for this child today. A tool that removes that judgement risks removing the part where the teacher is most valuable.

What the evidence actually shows

When I went back to check the claims I had been making, I found that "AI in education" is not one thing. A general chatbot, a carefully built tutor and an assistant sitting beside a human teacher are different interventions, tested on different learners against different comparisons with different measures. Four studies show why the design, the population and the outcome have to travel with every headline.

Four studies, four conditions

Different AI designs produce different outcomes, and different studies measure different things

A purpose-built tutor

[Kestin and colleagues, 2025](#) ↗

Peer-reviewed, Scientific Reports

194 Harvard undergraduates in physics, randomised across two lessons, against an active-learning classroom.

WHAT HAPPENED

Greater immediate gains in less time with a tutor designed for the course.

WHAT IT DOES NOT ESTABLISH

Retention, transfer, or that AI can replace teachers across a school education.

Practice up, learning down

[Bastani and colleagues, 2025](#) ↗

Peer-reviewed, PNAS

Nearly 1,000 high-school mathematics students in Turkey, randomised between no AI, a general chat interface and a tutor with safeguards.

WHAT HAPPENED

The general interface raised practice scores but the unaided exam fell by about 17%. Safeguards largely prevented the harm.

WHAT IT DOES NOT ESTABLISH

A positive effect on the unaided exam. Anything about tools not tested.

Access without engagement

[Oreopoulos and Low, 2026](#) ↗

NBER working paper, not peer reviewed

Middle-school mathematics across 18 Tennessee schools over two years, with and without the platform and tutor.

WHAT HAPPENED

Modest gains of roughly 0.06 to 0.08 standard deviations. Substantive tutor use was limited.

WHAT IT DOES NOT ESTABLISH

The incremental contribution of generative AI, since platform and tutor were bundled.

AI behind the human tutor

[Wang and colleagues, 2024, revised 2025](#) ↗

arXiv preprint

About 900 human tutors and 1,800 school students, randomised between tutors with and without AI guidance.

WHAT HAPPENED

Mastery rose by four percentage points when the tutor had AI guidance.

WHAT IT DOES NOT ESTABLISH

Effects of AI tutoring students directly, or outcomes beyond the session.

Immediate post-test gains, unaided performance, topic mastery and school-year achievement are not interchangeable outcomes.

Read together, they support something narrower and more useful than a verdict. A purpose-built tutor can improve immediate learning, at least in [a university physics course over two lessons](#). A general tool without safeguards can lift practice scores and hurt unaided performance. A platform-and-tutor bundle [made available to middle schools over two years](#) produced modest gains, while substantive tutor use remained limited, which suggests that engagement and implementation may matter as much as technical capability. [Giving AI guidance to a human tutor](#) rather than to the student raised the likelihood of topic mastery, which is a nudge towards strengthening human teaching rather than treating replacement as the default.

What they do not support is averaging. Immediate post-test gains, unaided performance, topic mastery and school-year achievement are not interchangeable outcomes, and a working paper, cited here from [the authors' abstract](#), is not a peer-

reviewed article. None of these studies validates a fully AI-native school, and none refutes one, because none of them tested one.

A LEDGER PRINCIPLE

The evidence supports disciplined experimentation, not a universal claim that AI works or that it does not.

Useful now, on one condition

I think we will go through a period in which AI sits on top of the system we have, so that students can make it through the system as it stands. That period is useful. Children facing this year's exams cannot wait for institutional redesign, and most of the practical good will be done here for some time.

The condition is that we do not confuse assisted performance with learning. The whole appeal of responsive help is that it can finally give the learner who needed one more example that example. The whole risk is that it gives the answer instead, and everyone, including the learner, mistakes the better answer for a better learner. The next essay takes that standard to a real product and asks what a tutor has to be built to do in order to meet it.

The measure of success is not how impressive the assisted answer becomes. It is how much more capable the learner becomes when the assistance is gone.

WRITTEN BY **Veemal Gungadin**

WITH ASSISTANCE FROM **Claude Fable 5.1** · Anthropic , **GPT-6 Astra** · OpenAI

Drafted and revised by the author. AI tools were used to sharpen structure and prose; conclusions are the author's own.

Published in *The Goodwill Ledger* at www.goodwill-led.com/essays/the-illusion-of-learning/