Schools Tighten AI Design as Homework AI Backfires and Structured Tools Advance

AI in EdTech Weekly

2026-06-22

Schools Tighten AI Design as Homework AI Backfires and Structured Tools Advance

By AI in EdTech Weekly • June 22, 2026

This week's coverage converged on a clear distinction: generic AI that completes work for students is undermining effort, while schools are investing in structured tutors, explicit AI-use rules, and bounded workflow tools. The result is a more practical, less hype-driven phase of AI adoption across K-12, higher ed, and self-directed learning.

The big development: education is drawing a harder line between answer machines and learning tools

This week's clearest signal was not a new model. It was a sharper distinction between AI that helps learners think and AI that helps them avoid thinking [1, 2, 3].

Ethan Mollick pointed to a recurring pattern: students naturally reach for AI on homework, but off-the-shelf chatbots act like assistants, not tutors, by providing answers that reduce mental effort and undermine learning [1]. He also cited a large study in China showing that when AI shortened homework time by lowering effort, test scores fell too [2].

> "Across studies, a theme: AI tutoring in support of classes is good, using AI to 'help' with homework is bad." [2]

The same tension is showing up in writing and assessment. One higher-ed analysis argued that AI has widened the college-readiness gap in writing by making it easier for students to produce polished text without doing the thinking writing is supposed to develop [4]. A narrower use on the Packback platform looked more promising: AI handled grammar and style feedback so instructors could focus on ideas, and student writing improved modestly over a semester [4].

EdSurge's podcast reached the same practical test for K-12: whether students are learning to think *with* AI or using it to bypass productive struggle [3].

Sam Altman said he expected schools to redesign quickly after ChatGPT, with projects that require AI but still stretch thinking, yet he still sees no significant systemic change across education 3.5 years later and warned that, without redesign, critical thinking skills could atrophy [5].

Stanford CS153 Frontier Systems | Scale, AGI, and the Future of Everything (29:04)

Mollick's practical response is more concrete: more in-class assignments, AI tutors that challenge rather than answer, and prompts that use AI as a critic during debate or argument rather than as a completion engine [6].

Even commentary on higher ed outcomes is moving in this direction. The AI in Education Podcast highlighted Berkeley research suggesting that A grades rose 30% in take-home writing- and coding-heavy courses after ChatGPT, while one computer science course's failure rate reportedly rose from 7% to 35% when students later faced exams without AI [7]. Researchers from Australia, New Zealand, and China are now explicitly arguing for Socratic AI companions that prompt reflection and understanding, rather than vanilla chatbots that act as a crutch [7].

The response is shifting from bans to explicit learning design

Higher ed and K-12 are starting to turn that insight into design choices instead of generic rules. Lance Eaton argues that year 5 of generative AI in higher education should be the year of program-level curriculum change, because students are still graduating after four years of scattershot experiences: one professor requires AI, another bans it, another treats it as misconduct, another never names it [8, 9]. His recommendation is not to simply teach the tool or ban it, but to map where students should first encounter AI, where they should use it, where they should work without it, and where they should learn refusal, verification, disclosure, and judgment [8].

A related UK discussion is pushing even further upstream. Rose Luckin argues that if AI can master knowledge-heavy curricula faster and more accurately than humans, schools need to shift toward richer human capabilities: creativity, problem-solving, metacognition, resilience, empathy, and better assessment of those skills [10].

K-12 policy is also getting more operational. Microsoft Teams now lets educators set assignment-level AI expectations: full AI use, editing only, brainstorming only, or no AI use, with customizable labels and defaults [11]. Students see the guideline when they open the assignment and, if their school enables it, a direct button to open Copilot [11].

Microsoft Teams Assignments: 6 New Features for Teachers (0:15)

That product choice matches district-level policy thinking. Tech & Learning highlighted a three-part framework for districts starting from scratch: understand how students and teachers are already using AI, protect student data and personally identifiable information, and address academic integrity without taking students out of the driver's seat of their own learning [12]. It also draws a distinction between prohibition-heavy acceptable use policies and responsible use policies that explain reasoning and treat students as participants [12]. In Massachusetts, Shrewsbury Public Schools built five pillars around student preparation, student learning tools, staff tools, guardrails, and academic integrity, aligned to the district's Portrait of a Graduate rather than to any single product [13].

This is also showing up in professional development. Hillsborough County Public Schools, a district serving more than 200,000 students, put 1,000 educators through a summer week of training on responsible use of MagicSchool AI [14]. The underlying message is similar to Monica Burns' advice: start AI guidance with the kind of thinking you want students to do, not compliance alone [15].

Tools are getting more specific — and their limits are clearer

The most credible product activity this week was not "AI does everything." It was AI being inserted into narrower learning workflows [11, 16].

In Microsoft's education stack, Learning Zone lets educators attach AI-built interactive lessons directly to Teams assignments, render them inside Teams for students, and provide built-in checks and feedback [11]. Those lessons can also draw from partner content including NASA, Figma, and Minecraft [11]. Rubric generation is becoming more constrained as well: when teachers create a rubric with AI, the standards already attached to the assignment are automatically carried into the rubric [11]. The limitation matters: AI lesson generation requires a Copilot Plus PC, even though students can complete the lesson in Teams on their side [11].

At the school operations level, the near-term gains remain mostly administrative. One principal described using AI daily to turn state memos into slide decks, teacher texts into parent messages, long emotional emails into summaries and replies, contract PDFs into queryable answers, and scattered event details into calendar entries — saving "a few hours" a day and freeing more time for students and staff [17]. Monica Burns describes the same division of labor more generally: AI can do the drafting, formatting, and structuring, but human review, edits, and knowledge of students still shape the final product [16].

For self-directed learning, platforms are getting more interactive. Copilot Notebooks is now available without paid student Copilot licenses, and one suggested use is uploading a curriculum to generate study guides, activities, and info-

graphics [7]. Google NotebookLM is already being used by students to turn class slides into self-quizzes and answer keys, reinforcing retrieval practice [18]. Andrew Ng is pushing toward a more conversational model in CodeDream.ai, where learners interact through simulated video calls and embedded JavaScript demos instead of passively watching static videos [19]. But Ng's verdict is restrained: online learning tools are better than they were 10 years ago, not yet truly transformed [19].

Adoption, meanwhile, remains a constraint. Mollick says AI interfaces like chatbots, Codex, and NotebookLM are not intuitive in practice and contain "a dozen little tricks and traps" that block effective use [20]. He also says many people never get past the difficult first hour, which keeps AI in the "kind of like Google" box [21]. Chalkbeat's reporting on a Stanford AI tutor study shows what that looks like in schools: human guidance increased use by only 1-4 minutes a week, many students never logged on, total time stayed far below the 30 minutes a week needed for reading gains, and there was no meaningful difference in reading scores [22].

> "The challenge isn't just building good AI tools. It's really getting students to use them, and that seems to take the same type of intentional design that we've learned matters with other ed tech interventions and tutoring." [22]

AI-native models are expanding — but not all in the same direction

Some schools are no longer treating AI as an add-on. They are designing schedules, staffing, and pedagogy around it from the start.

At Alpha School, students spend two morning hours on personalized academic work with AI tutors or adaptive apps that give immediate feedback and let students advance on mastery [23]. Afternoons shift to four hours of team-based life skills, projects, collaboration, and conversation [23]. Alpha argues that this split lets AI handle individualized cognitive work while freeing more time for authentic social learning, and says its classes rank in the top 1-2% nationally [23]. Reporting from Michael Horn's microschool series reinforces the broader pattern: the most interesting schools are not simply maximizing AI use; they are being explicit about where AI belongs and what human capabilities — autonomy, entrepreneurship, deep research, feedback, and strong foundations — they still want school to build [24].

Alpha is also trying to rebut a common critique directly. Its leaders say AI-first schooling should produce more thinking, not less, and are pairing the model with explicit humanities work, including students reading Tocqueville's *Democracy in America* and debating it for the age of AI as part of building "philosopher-builders" [25, 26].

Higher education is seeing its own AI-native experiments. The AI in Education

Podcast highlighted a new Italian online university built from scratch around AI optimization, serving 112,000 students with 400 academic staff [7]. That is a radically different staffing model from legacy higher ed. But Andrew Ng's comments are a useful counterweight: what people need to learn is changing quickly — coding agents, AI building blocks, and broader product skills — yet the delivery of training is still being reinvented in real time [19].

Economics are starting to shape product design

For edtech buyers and investors, one of the most practical notes this week came from e-Literate: current AI economics do not fit education's usual software model [27].

Schools and colleges budget around fixed annual costs, while metered AI introduces variable usage that can spike unpredictably [27]. e-Literate points to enterprise examples of blown token budgets, revoked licenses, and large unexpected bills as signs of what happens when usage limits are weak [27]. The likely consequence is product design, not just procurement, changing: mainstream platforms are more likely to ship constrained AI actions, narrow buttons, predefined workflows, and usage caps than open-ended magic text boxes or expensive multi-agent systems [27]. That slowdown may frustrate some vendors, but it could also reduce the odds that education scales the wrong tools too quickly [27].

What This Means

- **For K-12 leaders:** Redesigning homework and assessment is now harder to avoid. The evidence and commentary this week point toward more in-class checks, explicit AI-use expectations on assignments, and tutor-like AI that preserves effort instead of replacing it [2, 6, 11, 3].
- **For higher ed:** Course-by-course AI rules are too inconsistent. Program-level maps of where students should use, refuse, verify, and disclose AI are becoming a more practical governance model [9, 8].
- **For teachers and L&D teams:** The best short-term use cases remain bounded ones — interactive lesson building, standards-aligned rubrics, administrative drafting, and study supports — with human review still central [11, 17, 16].
- **For edtech builders and investors:** Capability is not enough. Products need engagement, usability, and cost discipline. Low-usage tutoring pilots and unsustainable token economics can kill otherwise promising ideas [22, 27].
- **For learners:** AI is most useful when it behaves more like a critic, coach, or quiz-maker than an answer machine [6, 7, 18].

Watch This Space

- **Socratic companions and public-interest tutors:** Researchers are pushing companion-style AI that promotes reflection, while Mollick argues universal tutors are now technically plausible if built with public R&D, transparency, and the right scaffolding [7, 28, 29].
- **Mainstream platforms embedding guardrails:** Assignment-level AI labels in Teams suggest more classroom software will make AI expectations visible inside the workflow, not just in policy documents [11].
- **Next-generation study platforms:** Khan Academy says its next launch will combine trusted content with AI tools to help students persist through hard learning, while Copilot Notebooks and CodeDream point to more interactive self-study formats [30, 7, 19].
- **AI-native school builders:** Alpha's summer internship and new engineering cohort show schools investing directly in building learning apps, not just buying them [31, 32].
- **Cost-shaped AI design:** Expect more predefined AI actions, fewer unlimited chat interfaces, and closer scrutiny of whether usage actually translates into learning gains [27, 22].

Sources

1. X post by @emollick
2. X post by @emollick
3. EdSurge Podcast: Your Kids Know More About AI Than You Do
4. Rethinking Academic Support for Today's Learners
5. Stanford CS153 Frontier Systems | Scale, AGI, and the Future of Everything
6. The AI Skills Nobody is Teaching (And Everyone Needs) | AI Expert Ethan Mollick
7. More A's, More Fails: What AI Is Really Doing to Student Performance
8. Agents, Agency, and Action…
9. substack
10. What if… we really wanted to prepare young people for the age of artificial intelligence?
11. Microsoft Teams Assignments: 6 New Features for Teachers
12. Bad AI Policy Is Worse Than No Policy at All. How to Build One That Works
13. Creating 5 Pillars To Guide AI Use In Your District
14. X post by @adeelorama
15. X post by @ClassTechTips
16. X post by @ClassTechTips
17. A Day In The Life Of A Principal Using AI
18. Stop Rereading Notes: How Students Actually Learn - Episode 943 with Blake Harvard

19. The Future of AI Agents with Andrew Ng | Interrupt 26
20. X post by @emollick
21. X post by @emollick
22. Research on AI tutoring ran into a problem: Most students wouldn't use it
23. The 2-Hour Cure for Your Kid's AI Loneliness
24. Tomorrow Schools: 7 Microschools Offering a Window Into the Future
25. X post by @jliemandt
26. X post by @CameronSorsby
27. Today's AI is Economically Unsustainable for Education
28. X post by @emollick
29. X post by @emollick
30. X post by @khanacademy
31. X post by @jliemandt
32. X post by @Austen