

Brain-First AI Becomes the Rule as Schools Ground Tools and Tighten Oversight

AI in EdTech Weekly

2026-05-11

Brain-First AI Becomes the Rule as Schools Ground Tools and Tighten Oversight

By AI in EdTech Weekly • May 11, 2026


The week’s biggest shift is clearer boundary-setting: evidence and practice increasingly favor AI that scaffolds thinking over AI that replaces it. Schools are pairing grounded tools with tighter usage rules, while privacy and procurement concerns are reshaping how fast adoption can move.

Brain-first AI becomes the default design rule

The most important development this week is a sharper distinction between AI that supports learning and AI that substitutes for it. Stanford’s SCALE Initiative reviewed 20 causal studies on AI in K-12 and found that AI improves performance while students have access to it, but those gains weaken or disappear on independent assessment [1]. In one randomized study, AI users led by 22 points on immediate recall and comprehension, but the gap fell to 6 points after three weeks and was no longer significant; on synthesis, evaluation, and application, the no-AI control group led at every timepoint [1]. Other studies in the same review found better essay scores without knowledge transfer [1], weaker memory-related brain activity and recall among ChatGPT users [1], and a shift from “write, reread, evaluate, revise” toward “ask AI, accept output, ask again” [1].

“Brain first. AI second.” [2]

That phrase matched a broader classroom consensus. The MIT Media Lab writing study discussed by Philip Seyfried and Vicki Davis found an advantage for students who started with their own thinking and brought in AI later, leading to a “yes-and” approach: keep writers’ notebooks, paper books, and partner talk, then use AI to accelerate thinking after the early work is done [3, 4].



Brain First, AI Second: Teaching Writing in the AI Era (1:00)

The strongest pro-AI evidence points in the same direction. Sal Khan described Khanmigo as a tutor that gives hints, examples, and nudges rather than direct answers [5]. He also said only about 10-15% of students can use AI constructively on their own; many need help learning what to ask and how to engage [6]. In an independent study in India, Khan Academy plus a human “lab in charge” produced roughly a half-standard-deviation gain in seven months versus Khan Academy alone [6]. Ethan Mollick noted that peer-reviewed meta-analyses still find positive effects of GenAI on learning, with the strongest evidence coming from randomized AI tutor interventions [7].

The practical design rules are getting clearer: protect the first attempt, force self-assessment before AI access, use AI to reduce extraneous load rather than remove productive struggle, and give teachers visibility into prompts so they can spot substitution early [1]. In AI-rich classrooms, that raises rather than lowers the value of teacher content knowledge, because someone still has to validate accuracy, sequence knowledge, interpret nuance, and design assessments that require judgment instead of surface correctness [8].

Schools are moving from generic chatbots to grounded workflows

Some of the clearest implementations this week came from schools that constrained AI tightly; even strong AI-school advocates are now separating struc-

tured mastery systems from open-ended chatbots [9, 10, 11]. At Oran Park Anglican College, teachers built NotebookLM “brains” around curriculum documents, syllabi, universal design for learning principles, and writing scaffolds so outputs stay aligned with pedagogy and inclusion goals [12]. The school says AI use is saving about 52 hours per week on average while improving pedagogy, and every tool goes through staff training and a six-month pilot before wider rollout [12].



What we learned from teachers at Sydney’s Day of AI (1:20)

Student use is being structured just as deliberately. Oran Park launched a simple assessment scale — no AI, AI as assistant, AI as collaborator — so teachers can decide when and how AI is allowed [12]. At Roseville College, surveys found that 90% of students had used large language models and 56% were using them at least weekly, but only 23% had seen teachers model strong use [12]. That gap is pushing schools toward process-based assessment: one geography task now locks students’ research folders before class, lets them use AI to prepare materials, and then requires in-class work that reveals whether they actually engaged with the content [12]. A similar framing is emerging in policy design: once AI is allowed, the more useful question becomes whether a student engaged with it maturely — with enough effort, thinking, and persistence — not whether every imperfect interaction is an ethical failing [13].

Other classroom practices follow the same logic. Teacher Nathan Jones encourages AI for scaffolds, checklists, and task breakdowns, asks students to attach

prompt appendices, and explicitly teaches verification because hallucinations still happen [12]. That is one reason AI detectors are losing ground: teachers are shifting toward monitored class-time use and process conversations rather than trying to infer authorship after the fact [3]. Wayground’s new accommodations feature lets one staff member enter IEP-based accommodations once and apply them automatically across all of a student’s activities, reducing teacher overhead while keeping support consistent [14].

On the product side, the most useful capabilities are increasingly source-grounded and editable rather than one-shot magic. NotebookLM’s new auto-labeling organizes large notebooks into overlapping categories and lets users focus chats or generated artifacts on selected source groups [15, 16]. Gemini Canvas lets teachers build custom interactive tools without manual coding, but the educator stays “in the lead” and iterates the result [17]. Microsoft’s Teach module can generate standards-aligned Minecraft Education lesson plans, yet the output is still a draft that educators edit, enhance, and save into their own workflow [18]. Ethan Mollick’s warning on an AI-built physics simulation captures the limitation: even when something looks good after cursory checks, it still needs deeper verification before being used to teach students [19, 20].

AI fluency is becoming a core literacy

Higher education and K-12 are both moving beyond “can students use the tool?” toward “do they understand what the tool is doing?” Agnes Scott College will embed a three-part AI curriculum in the first-year experience starting in fall 2026, treating AI fluency as a core literacy alongside writing and quantitative reasoning [21]. Its distinction is useful: competency is basic operation; fluency means understanding limits, bias, ethics, and who benefits or bears the costs when AI is deployed [21]. The school argues that practical skills matter, but not at the expense of judgment, critical analysis, and accountability [21].

That broader literacy need is also showing up in K-12. At Philadelphia’s Marian Anderson Neighborhood Academy, middle schoolers researched AI’s effects on education, government, creativity, and the environment. Students described both upside — such as writing help or creative experimentation in Roblox and video editing — and downside, including cheating and the sense that AI is now hard to avoid because search engines surface AI overviews by default [22]. School leaders framed the work as an ongoing dialogue with families, educators, and public officials, not a one-off tech lesson [22].

Institutions are starting to back that shift with money and policy. The U.S. Department of Education is giving preference in discretionary grants to proposals that expand AI literacy, support ethical use, and improve student outcomes [1]. Google is putting \$10 million into AI skills across Asia-Pacific, explicitly centering teachers as the people who translate tools into classroom impact [1].

Adoption is colliding with privacy, trust, and procurement

The strongest policy signal came from New York City. The city recently canceled plans for a selective AI-focused high school after community criticism of both its admissions process and AI use in the classroom [23]. After that proposal was nixed, more than 100 New Yorkers demanded an AI moratorium at a marathon board meeting, echoing broader concerns about transparency, safety, and oversight [24, 1].

Those concerns are not happening in a vacuum. NYC’s new AI framework could expand AI use for tasks like lesson planning and translating materials for bilingual learners, but a state audit found the district lacks a complete inventory of third-party software and has already faced multiple breaches, including PowerSchool and Illuminate incidents [25]. Critics argue the framework raises privacy risks further and that the city should strengthen in-house infrastructure and oversight before accelerating AI adoption [25].

At the state level, the conversation is also getting more concrete. Vermont’s proposed H.650 would require edtech providers to register and be reviewed on design features including AI, geotracking, and targeted advertising [26]. Rhode Island’s Safe School Technology Act would restrict providers from activating audio or video functions outside school activities and ban use of location data [26]. As Andrew Marcinek argued, the answer is not simply to ban technology, but to build more intentional programs and communicate more clearly with parents [27]. Employers including Microsoft and Anthropic are already signaling demand for workers who understand AI alongside strong soft skills [27].

What This Means

- **For classroom design:** Build assignments so students think before AI enters the room. Independent brainstorming, first drafts, or problem attempts followed by AI critique fits the strongest evidence better than AI-first generation [1, 3, 5].
- **For school leaders:** The safer bet is not “allow AI” or “ban AI,” but structured use: assessment scales, prompt visibility, teacher modeling, and pilot-based rollout [12, 1].
- **For higher ed and L&D teams:** AI training should move beyond prompt tips toward fluency — limits, bias, evaluation, and when not to delegate the work [21].
- **For product teams and buyers:** Grounding matters. Tools tied to curriculum docs, standards, or user-provided sources are showing clearer value than open-ended chat, but they still need expert review for accuracy and alignment [12, 15, 18, 20].
- **For policymakers and families:** Privacy and trust are now adoption constraints, not side issues. If districts cannot account for third-party tools or data exposure, AI expansion will keep meeting resistance [25].
- **For self-directed learners:** AI can help organize complex projects and

materials, but tools that remove all friction may also remove some of the struggle that produces learning [28, 29].

Watch This Space

- **Source-grounded personal study spaces:** Gemini Notebooks are being positioned as a place to gather drafts, requirements, and deadlines in one AI-assisted workspace [28, 30]. Lance Eaton describes a similar agentic pattern: semantically analyzing 200+ course materials, tagging them, mapping connections, and building a custom reader with recommendations — useful for reducing organizational friction, but still in tension with learning’s need for deliberate friction [29].
- **Young builders using AI for real work:** OpenAI’s ChatGPT Futures highlighted students using AI to map 1.5 million previously unknown objects in space, make 100 million galaxy images searchable, detect disaster survivors through walls and debris, preserve endangered languages, and reroute more than 5 million pounds of unsold inventory from landfills [31]. In Philly, middle schoolers are already using AI for Roblox game coding and video editing outside school [22].
- **Rules for AI in schools:** Watch for more states to move beyond generic AI statements and toward product-level requirements covering design features, data access, and device permissions [26].

Sources

1. The Cognitive Debt Problem
2. X post by @coolcatteacher
3. Brain First, AI Second: Teaching Writing in the AI Era
4. X post by @coolcatteacher
5. Scaling Education: Access, AI, and Impact – Conversation between Sal Khan and Sajida Shaikh
6. The Future of Teaching & Learning: Education in the AI Era | 2026 Common Sense Summit
7. X post by @emollick
8. In An AI Classroom, Content Knowledge Matters More Than Ever
9. Rethinking Education & The 2-Hour School Day That’s Changing Everything w/ MacKenzie Price
10. “Why our family loves Alpha School”
11. AI Schools Are Here: How kids learn 2h/day and become top 1% nationally | MacKenzie Price
12. What we learned from teachers at Sydney’s Day of AI
13. Mature AI Use vs. Immature AI Use
14. X post by @KyleNiemis
15. X post by @stevenbjohnson
16. X post by @NotebookLM

17. Vibe Coding with Gemini Canvas
18. How to use the new Minecraft Education lesson plan generation tool in Microsoft 365
19. X post by @emollick
20. X post by @emollick
21. What Colleges Get Wrong About AI Education
22. Philly middle schoolers are examining AI — and questioning its impact on their lives
23. 5 new public schools opening this fall in the Bronx and Queens
24. Pennsylvania overhauled its graduation requirements. Are graduating high schoolers any better off?
25. State audit slams NYC schools for lack of student data privacy oversight
26. Screen Time Concerns Lead to Backlash Against Edtech Vetting Process
27. Edtech's Big Tobacco Moment Is Here. Schools Can't Afford to Miss the AI Reckoning That Follows
28. X post by @GeminiApp
29. Exploring Agentic AI
30. X post by @GeminiApp
31. X post by @OpenAI