

# AI in Learning Turns Toward Guided Practice as Schools Scale and Safety Rules Tighten

AI in EdTech Weekly

2026-03-09

## AI in Learning Turns Toward Guided Practice as Schools Scale and Safety Rules Tighten

*By AI in EdTech Weekly • March 9, 2026*

This week’s strongest signal is a design shift: the most credible uses of AI in learning are moving away from instant answers and toward guided practice, teacher control, and clearer guardrails. The brief also covers new school and university deployments, rising workforce pressure for AI competence, new tools such as NotebookLM video overviews, and the policies now shaping student-facing AI.

### **The lead — AI is being redesigned around guided practice, not just answers**

Several sources this week converged on the same design principle: AI helps learning most when it keeps the learner thinking, gives teachers more control, and fits inside a human system of practice and accountability [1, 2, 3].

Chris Yue Fu’s eight-week study of 15 undergraduates found students often used AI summaries as the thing they read, not as support for reading, and only 4.3% of their prompts used effective strategies even after instruction [1]. Fu’s takeaway was not to remove AI from reading, but to redesign it so teachers can set goals in advance and the system pushes students into higher-order questions instead of ending the task after one answer [1].

“When a student asks for a summary and gets one, the system has done its job, but the student hasn’t done theirs.” [1]

The same pattern showed up elsewhere. Ethan Mollick argues that learners gain when AI supports coding, but not when it replaces the intellectual work; he also pointed to findings that *vibecoding* can hurt developers’ ability to read, write, debug, and understand code, without producing a statistically significant speed

gain [4, 5]. Justin Reich’s tutoring work frames the distinction similarly: good tutors do not just answer questions; they **question answers**, and schools still need small experiments rather than sweeping assumptions about best practice [3]. One cited high-school math study made the tradeoff stark: ChatGPT access increased correct practice problems by 48% but lowered actual test performance by 17% [6].

Some schools are already building around that insight. In Italy’s GEMI project, teachers who began skeptical of Gemini ended up using it as a thinking partner, including generating deliberate errors in a literature text for students to detect, while keeping the educational relationship central and using NotebookLM to support students with specific learning needs [7].



*Learn how schools in Italy are using Gemini to help students improve their critical thinking (2:12)*

## **Theme 2 — AI is becoming part of the school operating model**

The most concrete AI-school expansion this week came from Alpha School, which is opening a new K-8 campus in The Woodlands in fall 2026. The operational detail is notable: leadership says demand is not the main growth constraint; real estate is. Alpha also added Nate Eliason to expand AI and entrepreneurship at the high-school level [8, 9, 10].

Elsewhere, the infrastructure is getting more local. A school in China reportedly repurposed M1 Ultra Macs, clustered them with Exo, ingested its full school corpus, and gave each student and teacher a personalized, free, private AI agent grounded in real school data [11]. Alpha leadership highlighted that example as part of global benchmarking for what school AI may look like outside the U.S. mainstream [12].

Khan Academy’s recent signals show how much implementation still matters. Khanmigo is framed as immediate-feedback tutoring in math and writing, but Sal Khan says the learning gains come from more practice at a student’s learning edge with teachers in the loop, not from AI alone [2]. Reported thresholds of roughly 18 hours a year or 60 skills to proficient are associated with meaningful gains, an India study found a 0.44 effect size at reasonable dosage, and one Newark district area reported twice average state-test growth [2]. The district playbook is decidedly human: training, leadership, support for teachers, and soft accountability rather than punitive mandates, with support priced at about \$15 per student per year [2].

### **Theme 3 — AI competence is moving into curricula, credentials, and workforce training**

Higher education is moving from generic AI talk to domain-specific expectations. Purdue now requires AI competency for all graduates, with each discipline defining what job-ready use looks like [13]. The University of Sydney and UTS partnered with Harvey AI to prepare law students for a legal AI system already used in professional practice [13]. The University of Manchester says it is rolling out Microsoft 365 Copilot and training to all 65,000 staff and students [13].

Self-directed and professional learning are following the same path. DeepLearning.AI released a free AI Skill Builder to help learners assess what to learn next, alongside a new JAX course on building and training a 20M-parameter MiniGPT-style model and a broader roadmap centered on agents, external data, evaluation loops, alignment with human intent, and interaction with tools [14, 15, 16]. EMERGai opened applications for an NSF-funded institute that gives early- and mid-career STEM education researchers at U.S. resource-limited institutions stipends, training, and support to use GenAI ethically across literature review, data collection, analysis, interpretation, and writing [17].

Corporate training is scaling too. Gauntlet AI says it has worked with more than 80 training and hiring partners in its first year and expects to more than double that figure, while also finding that bringing product managers alongside engineers matters because the shift is cultural and tooling-related, not just technical [18, 19]. The labor-market pressure behind these moves is showing up in anecdotes from computing education as well: in one account of a Berkeley CS cohort, 31 of 340 majors had offers, while postings increasingly asked for AI/ML experience and the ability to review AI-generated code [20].

## **Theme 4 — Guardrails are becoming a first-order product requirement**

Child safety and governance are moving from side discussions to product requirements. Under the UK Online Safety Act, Ofcom says services must use age assurance for harmful content, prevent algorithms from recommending that content to children, and carry out child risk assessments when they launch generative AI features. Regulators also say some services should simply not be available to children [21].

Experts in the same discussion flagged newer harms that are harder to regulate cleanly: emotional dependency on chatbots, harmful advice, deepfakes, explicit chatbot conversations, bias, and personal-data exposure. Their recommendation was that schools discuss AI ethics and safeguarding explicitly and bring parents into those conversations, especially since AI was the top issue in one survey of 800 schools about online-safety conversations with families [21].

That regulatory lens is colliding with product reality. Google’s Gemini API terms say developers must not use the service in apps directed toward or likely to be accessed by people under 18 [13]. In student mental health, Alongside is now used in more than 200 U.S. schools and costs about \$10 per student per year; one Florida counselor credited it with surfacing a severe self-harm alert and helping handle routine problems so human staff could focus on crises [22]. But clinicians and researchers cited by EdSurge warn that AI lacks human discernment, should not substitute for counseling, and can encourage parasocial attachment if it signals emotional reciprocity [22].

In higher ed, adoption is already ahead of policy. A survey discussed on the AI in Education Podcast found 73% of respondents using AI daily or weekly, more than half using tools not provided by their institution, and only 13% saying their university measures ROI [13]. A complementary argument from edtech researchers is that AI tools should be judged on efficacy, effectiveness, equity, ethics, and environment — not adoption alone [23].

## **Theme 5 — The tool layer keeps expanding, but capability and limitation are arriving together**

NotebookLM’s biggest education-facing release this week was Cinematic Video Overviews. Gemini chooses a format and visual style, critiques its own footage, and turns source material into bespoke videos from a user’s sources [24, 25]. The limitation is equally clear: the feature is fully rolled out only to Ultra users in English for now, with Pro users still waiting [24, 26].

Groovelit shows the opposite end of the stack: a free grades 4-10 writing platform where students write in timed rounds and get live AI feedback on grammar, relevance, vocabulary, and engagement. Teachers can align prompts to curriculum, review aggregated data, and support English language learners with adjusted difficulty [27].

At the curriculum-engineering layer, Austen Allred says he is testing AI against 21 learning-science requirements, including spaced repetition, retrieval practice, semantic tree traversal, and mastery-based progression [28]. His own limitation note is blunt: the software harness was easy; getting the AI to follow it reliably — and stop inventing fake reviews or user numbers — was not [28, 29].

## What This Means

- **For K-12 leaders:** prioritize tools that preserve productive struggle. The consistent pattern across reading, tutoring, coding, and classroom examples is that AI is strongest when it asks better questions, provides feedback, or scaffolds practice — not when it substitutes for reading, writing, or reasoning [1, 4, 3, 6].
- **For school systems and edtech operators:** treat AI adoption as an operating-model problem, not just a product choice. Reported gains around Khanmigo depend on training, leadership, dosage, and teacher engagement, while Alpha’s next bottleneck is physical expansion capacity, not interest [2, 9].
- **For higher ed and workforce programs:** move from generic AI literacy to domain-specific workflows. Purdue, Harvey AI in law, Manchester’s Copilot rollout, and Gauntlet’s mixed PM/engineer cohorts all point to AI competence becoming contextual, team-based, and tied to real tools learners will encounter at work [13, 19].
- **For edtech buyers and investors:** under-18 access rules, child-risk design, and measurable impact are moving to the center of procurement. Gemini’s age terms, Ofcom’s expectations, Alongside’s limits in mental-health use, and the low rate of ROI measurement in higher ed all point to a more demanding buying environment [13, 21, 22].
- **For self-directed learners:** build sooner, but keep ownership of the thinking. DeepLearning.AI warns against staying in tutorial mode, and multiple sources this week warned that outsourcing cognition to AI weakens learning even when it makes the task feel easier [30, 14, 4, 31].

## Watch This Space

- **Teacher-controlled tutors that keep the conversation going:** especially tools that ask follow-up questions, set teacher goals, and steer students toward reasoning rather than one-shot answers [1, 3].
- **Whether private, local school agents grounded in institutional data move beyond isolated examples:** the China deployment is a concrete model to monitor [11].
- **Whether discipline-specific AI requirements move beyond the current university rollouts:** this is already visible in graduation requirements, campus-wide copilots, and workplace-linked tools [13].
- **How youth-facing companions and mental-health-adjacent bots are governed:** one reported figure says 72% of teens have used AI for

companionship at least once and 52% do so daily, even as regulators and clinicians flag emotional dependency as an emerging harm [6, 21, 22].

- **How multimodal study aids are used in practice:** from NotebookLM video generation to dual-voice podcasts for text comprehension, AI is expanding how source material gets remixed for learners [24, 7].

---

## Sources

1. The Right AI Can Help Students With Assigned Readings, Suggests New Research
2. Fireside Chat with Sal Khan and Asst. Superintendent Phil Misecko
3. EduTrends ep. 79 - Overreliance on AI for Thinking and Learning with Justin Reich
4. X post by @emollick
5. X post by @aarondotdev
6. ASK ME ANYTHING #27: "Is AI rotting our kids' brains?"
7. Learn how schools in Italy are using Gemini to help students improve their critical thinking
8. X post by @HelloWoodlands
9. X post by @jliemandt
10. X post by @nateliason
11. X post by @alexocheema
12. X post by @jliemandt
13. From Classrooms to Careers: The New AI Skills Race
14. X post by @DeepLearningAI
15. X post by @DeepLearningAI
16. X post by @DeepLearningAI
17. Professional development institute applications open: Expanding the Methods of Education Research with Generative AI (EMERGai)
18. X post by @Austen
19. X post by @Austen
20. X post by @TechLayoffLover
21. 318 The Future of Child Online Safety: Insights from Ofcom LGFL
22. With Teens Comfortable Confiding in AI, Should Schools Embrace It for Mental Health Care?
23. The Learning Impact of AI Can and Must Be Benchmarked
24. X post by @NotebookLM
25. X post by @NotebookLM
26. X post by @NotebookLM
27. What is Groovelit and How Can I Use It To Teach Writing?
28. X post by @Austen
29. X post by @Austen
30. X post by @DeepLearningAI
31. X post by @MLStreetTalk