

Generative AI and Cognitive Stunting: Why Education Must Redesign Learning

The Economy Research Editorial^{1,2}

¹The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

²Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

This paper addresses the claim that children's increasing reliance on generative AI threatens a kind of cognitive atrophy. It shows the concern to be justified only if taken not as an ethical lapse by children, but as an institutional crisis in educational assessment practices, curricular progression and governance structures, designed for an earlier technological environment. Generative AI differs from search in that it does more than seek to reduce the cost of retrieval: it can also generate synthesis, explanation, argumentation, style and what appears to be understanding. It alters the way cognitive effort is distributed across learners, intermediary tools and educational institutions. This paper develops the distinction between AI as an effort reducer and AI as a systematic enabler of learning, focused on the varying impact of adoption, cognitive offloading, AI tutoring systems, assessment policy, classroom scaffolding and exposure to labor markets. It concludes that the right policy is neither prohibition nor passive acceptance, but a restructuring of the curriculum to focus on developmentally-balanced progression, tightening of formative assessment, clear institutional rules for AI use and protecting disadvantaged pupils from unmanaged AI exposure. The key challenge is how to best enable children to operate purposefully and freely with powerful new tools without sacrificing their cognitive skills.

1 Introduction - Generative AI and the Redistribution of Cognitive Work

The recent proposal to consider children's overdependence on generative AI as a form of "cognitive stunting"^[1] captures something significant, but not enough. The concern is that learners using new tools in place of effortful mental practice could weaken the cognitive muscles that education is supposed to train. And yet that metaphor can become shallow if used as a diagnosis of children rather than as a warning about institutional nostalgia. The policy problem is larger and more structural. Generative AI is not just another research tool and the question is not whether it constitutes "real intelligence." The more important observation is that it shifts the distribution of cognitive work in a number of ways: retrieval, explanation, summary, composition, drafting, editing and simulation of conversation, in the entire learning suite of activities.^[2] All of this simultaneously changes the motivations of children, students, teachers, administrators and governments. The key point is that the problem is not simply underuse of the mind, but a rapid change in the environment for which schools' assessment models, curricular logics and governance mechanisms were not designed.

That mismatch is even more urgent in 2023-26, because three developments have come together: adoption is exploding; the education system entered the AI era from a position of weakness; and AI's labor market relevance is no longer speculation. Adoption has proliferated and accelerated at a pace unlike any other technological innovation. In the United Kingdom, HEPI survey data reveal that student AI use is rising from 66% in 2024 to 92% in 2025^[3] and 95% in 2026.^[4] Pew data show the percentage of teens using ChatGPT for coursework doubling from 13% in 2023 to 26% in 2024, in the United States.^[5] The education system entered the AI era from a position of weakness rather than preparedness. According to the World Bank and UNESCO, the global learning crisis persists,^[6] with learning poverty very high in middle- and low-income countries and OECD PISA 2022 results record drops in mathematics and reading across OECD systems.^[7] The labor-market importance of AI is no longer based on guesswork: the IMF estimates that 60% of jobs in advanced economies are vulnerable to AI^[8] and the OECD concludes that generative AI is likely to accentuate regional gaps unless digital infrastructure, human capital training and public policy respond at pace.^[9] These three factors cannot be reconciled with treating AI as a pedagogical side-show: it is a permanent structural factor in human capital formation.

The argument therefore runs in two directions. The idea of generative AI as "Google 2.0" is useful only if it is understood carefully; it is a step further along a longer road of cognitive offloading through search, calculator and digital navigation. The difference is not metaphysical: it depends on what cognitive layer we offload, at what age, for what institutional benefits. Second, the most problematic action is not to try to turn back the historical clock and reintroduce AI into classrooms. Education was unwell in many ways before ChatGPT and these current failures contributed to why AI is such a disruptive innovation. When assessment rewards polished output over demonstrated understanding, students will turn on their own a tool that can produce quick answers. In this way, AI is exposing an educational landscape already weakened by its own failures: weak foundations, fragile assessments and deficient assessment feedback loops. The solution cannot be to deny that AI is here: the solution must be to fundamentally reshape education so that it can stand up to the powers of AI and to

integrate it where it strengthens learning and constrain it where it displaces the effort through which learning is formed.

2 Beyond “Google 2.0”: What Generative AI Actually Offloads

The comparison with Google is not trivial. Search engines have already normalized a key mode of cognitive offloading.^[10] Research on the “Google effect”^[11] in the literature recently concluded that internet searches influence memory, especially by transferring aspects of content to aspects of form that are more learnable. Similarly, a 2024 cognitive science model of offloading has argued that mental effort has value when compared to the light external work of reminders and that we offload more when internal maintenance is more burdensome. In this small way, generative AI now fits into a long line of tools that free up mental resources by minimizing demands on memory, directly or indirectly. It is not the first technology to do so and it is not the first technology that maximizes retrieval at the expense of remembering. This bit of continuity in our history is helpful because it prevents apocalyptic moral panic. Our cognition has always been distributed across media, artifacts, institutions and other tools.^[12]

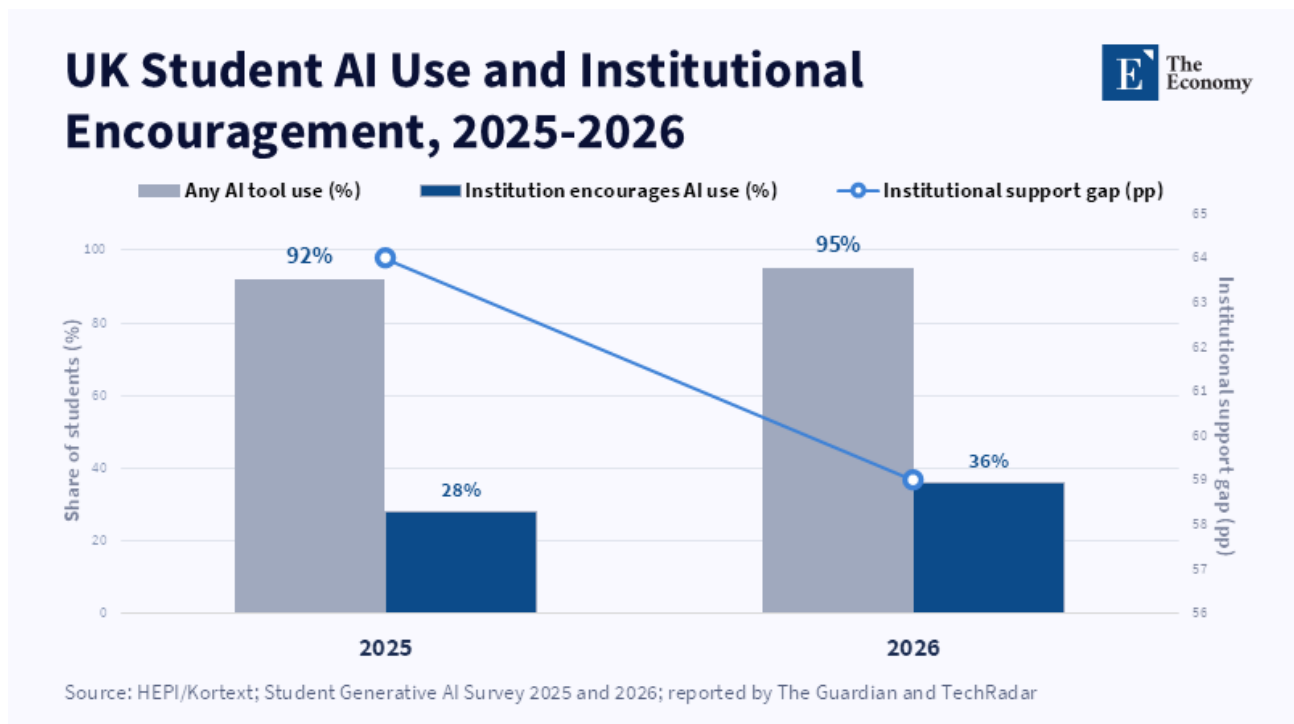


Figure 1: AI use among UK students is now near-universal, while institutional encouragement still lags far behind adoption.

But the continuity can be misleading, since search will mainly externalize retrieval while leaving interpretation to the student. The learner still needed to compare sources, then figure out what each source meant and how it went together and then convert the source material into a final argument or solution to produce a persuasive final answer. Generative AI pushes all those intermediate processes together into one slick-looking output, at once. It can do more than this to support the student in finding the facts. It can also provide a rationale, a draft, a tone, a structure and increasingly a dialogue partner she can converse with in refining the resulting product to a point where it looks finished. And this impact is becoming obvious in the ways students

move around sources: Pew found that on encountering AI-summarized pages, Google searchers clicked on standard search links only 8% of the time,^[13] from a higher baseline of 15% for unaffected pages,^[14] while clicks on links within the summary itself only occurred 1% of the time;^[15] and at the same time, general language models can still hallucinate facts and cite non-existent sources,^[16] as evident in Nature’s efforts in hallucinometric detection and OpenAI’s own acknowledgement that its system for detecting AI-written text was withdrawn for unreliability^[17] and that ChatGPT cannot reliably identify AI-used language. The point is not that AI responses are too hard to use. It is that the interface encouraged source blindness even in the presence of a need for confirmation.




Figure 2: AI summaries appear to reduce source-checking by keeping users inside the answer interface.

For education, that difference is significant. Search technology reduced the retrieval cost of information; generative AI reduces the cost of appearing to understand. These are not equivalent in education. If the student is still going to produce the argument, then educational value can endure considerable reduction in search costs. But if the tool produces the argument itself, then the important issue is whether the student is, in the process of preparing for applying knowledge, still doing the generative, effortful work of developing durable schemas. The supervision task has also shifted. Using web searches, learners compare multiple external objects; using generative AI, learners compare one internally generated answer whose fluency might conceal error and whose style might create an illusion of mastery. One study of 319 knowledge workers by Microsoft researchers reported that increased confidence in generative AI was associated with less critical thinking,^[18] just as increased self-confidence was associated with more critical thinking. In this way, the burden of judgment is moved away from information search and toward the authority we give to our own judgment. This is an extremely difficult challenge for novices but more manageable for experts.

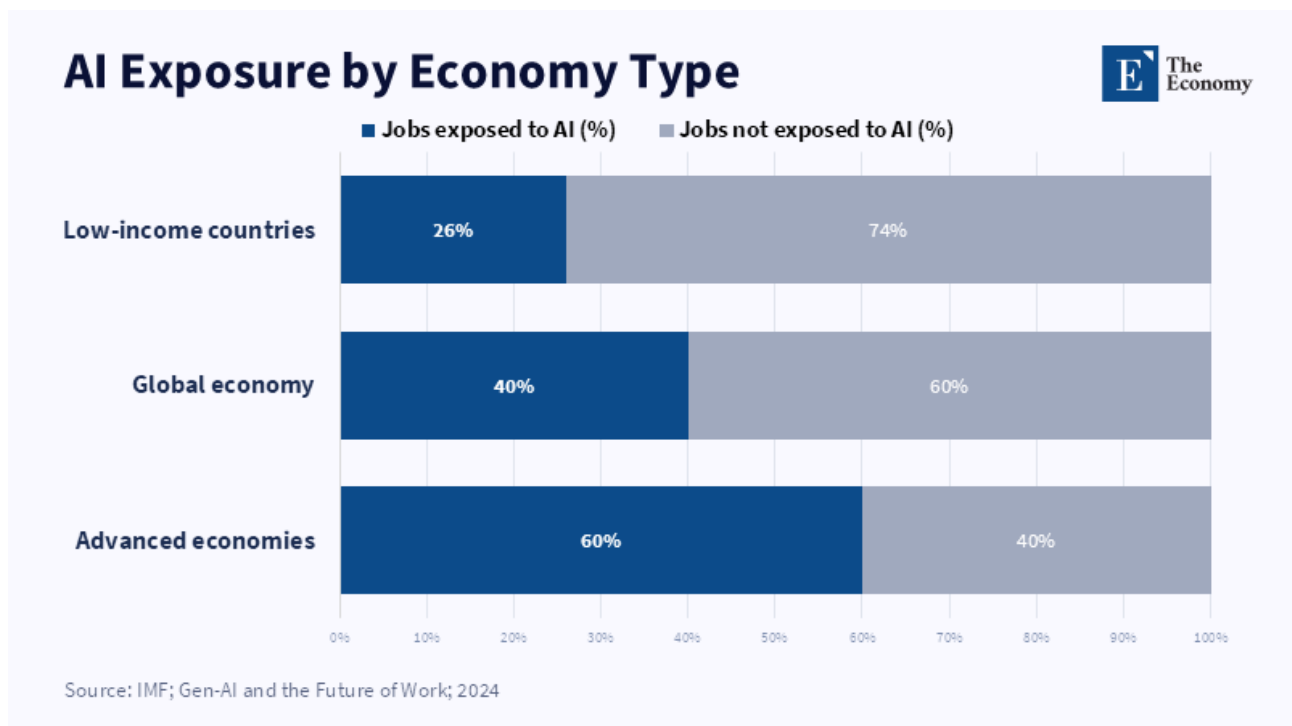


Figure 3: AI exposure is highest in advanced economies, where knowledge-intensive work is more concentrated.

This is why analogies to Gutenberg, the phone, the calculator or cloud computing are useful only up to a point. Every new medium of communication has been hounded as the corrupting force, undercutting skills that came before; some of those fears were mistaken, yet no medium intervenes at the same node of human action, all at the same depth. Print expanded access to the message, yet still needed to be read. The phone’s instant transmission of speech upended the letter of recommendation, seminar document and statistical metanarrative, yet lacked the capacity to produce those messages in a speaker’s script. The calculator hastened the solution, but was dumb when it came to formulating the Pythagorean theorem, a theoretical explanation, or a persuasive answer in an expository essay: generative AI approaches a general-purpose synthesis. It already descends further down the sequence from human problem to human answer than search ever did. The question is therefore not whether it is “just search.” The question is whether educators can recognize the offloading form that frees cognition for tomorrow’s learning, from that which displaces today’s deep effort.

3 Cognitive Stunting as a Developmental and Institutional Risk

The argument about cognitive stunting is most compelling when applied to children and novices, for it is a misconception to view learning as the simple gathering of correct associations. It ignores how fundamental, long-term, repetition-and-correction, re-application and self-explanation are and how slow conceptual development can be. The recent literature on desirable difficulties^[19] continues to highlight how the worst forms of effort, though so maligned in the moment, are actually responsible for the gains in long-term retention and transfer we seek from any difficult task. So too, then, is the educational danger ahead with generative AI: not just academic cheating, but the risk that a task-optimization, performance-oriented, generative AI will remove too many of the conceptual frictions through which understanding is built, thus risking learning itself. The Brookings argument

is strongest when understood in this narrower developmental sense, when they seek to make: some artifacts of development are simply not next in line for direct migration. UNESCO’s less technical preview of a computer from the perspective of teachers echoes the same insight from the opposite direction: that digital activity is meant to serve as a context-augmenter for human work, not a replacement for it.^[20]




Figure 4: Mathematics recorded the steepest OECD-average decline before schools fully entered the generative-AI era.

The bulk of modern data from school-age participants is illuminating the process. Pew research in 2023 shows nearly 25% of the U.S. teens using ChatGPT for schoolwork in 2024, up significantly from the prior year and twice as popular among older teens. An experimental study carried out in 2024 with 106 secondary-school students from Germany, though, showed a perceived usefulness of chatbots despite limited prior experience with them, combined with significant underuse of adaptation features that offer the greatest pedagogical value without active instructor demand. The experimental subsample that received a brief basic instruction on how to listen for cues requesting simplification, summaries, or mid-level writing made twice as many adaptation requests as the untrained group.^[21] No difference existed in induced knowledge gains. The pattern is significant here. Perceived usefulness is a matter of apparent superficial affordance with systems that have long-term pedagogical impact that must be discerned.

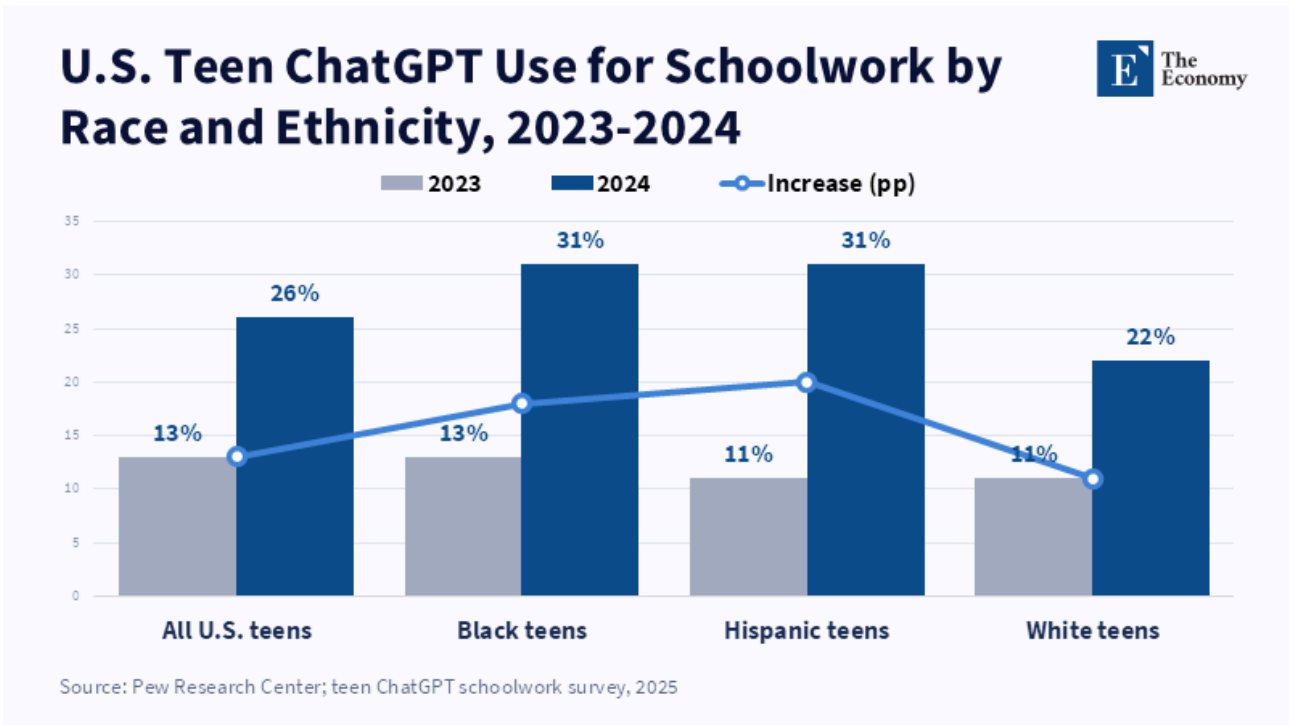


Figure 5: ChatGPT use for schoolwork doubled among U.S. teens overall, with especially sharp increases among Black and Hispanic teens.

The pattern reappears under more demanding conditions. A meta-analysis of a trio of articles from the Harvard Educational Review on college programming courses (which were reviewed by Stanford) identified two diametrically opposed influences of the use of LLMs. Students performed better on automated test assessments when they relied on the models as conversational tutors, prompting them to seek explanations and forcing them to draw inferences with the material. Students performed worse when they relied on the models to generate solutions on practice problems,^[22] delaying the effort they would otherwise have to put forth in the question. Learners overrate the mental improvements gained vis-à-vis external aids, relative to the improvement they make internally on their own, confirming prior results that subjective claims of benefit outstrip actual effects.^[23] It is no shock that a learner would double down on that illusion: the effectiveness monitor of the self is itself evolving. Once students are able to merge the tentative responsiveness of fluid completion with the appearance of expertise, the ability to police the gap between their useful self and their performance, aided by the mask, evaporates.

The "Your Brain on ChatGPT" project has intensified this concern,^[24] although its findings should be treated as preliminary for precisely this reason. In a small and certainly preliminary longitudinal essay writing experiment between an LLM condition, a search condition and a no-tools condition, the researchers observed attenuated neural activation, aversions to the produced essay and poorer retention in heavy LLM users^[25] relative to search users and uninformed essay-writers. Caution is necessary in regard to overgeneralizing from this modest sample size and in regard to taking its conclusions as absolutes, but the comparison of the search grouping and the LLM group should be informative for the argument made here. Not all forms of digital assistance appear to result in the same loss of the active selection and binding processes that produce authorship. Search preserves a greater portion of the process than automated text composition. That does not mean that

search is pedagogically neutral, but it does suggest that concurrent AI and instant and novel text production create new offloading risks.




Figure 6: Older students showed the largest decline in time spent on AI-susceptible math problems, suggesting higher outsourcing risk where tasks are easier to automate.

At this point, the focus must also be moved away from children as moral subjects and towards incentive systems as the institutions that operate them. The reason essay-writing has become vulnerable to deception is obvious: it is, in testing conditions, comparatively rational for a student to have the work done for her if it is possible. If the essay is treated only as a finished artifact one totemic essay as a piece of a project is merely a finished piece and if a user of a chatbot has a few simple commands, it would not be very difficult for her to produce the former without doing the latter. But the nature of assessment practice is misaligned. OpenAI’s own abandoned classifier is once again illustrative: OpenAI announced that its classifier differentiated at a mere 26% for challenge set samples and signaled human authorship at a rate of 9 percent, all in a tool that OpenAI subsequently abandoned because it was too inaccurate. Detection bureaucracies to the uninformed and punitive will not last in credibility: they evoke false accusations under conditions where the incentives, like much of the education system, are still fundamentally misaligned.

Distributional consequences are also severe. RAND’s large nationally representative survey of teachers conducted in 2024 found that approximately 25% of all U.S. teachers of core subjects utilized AI tools as scaffolds to practice in lesson planning and instruction, but only 18% of principals indicated that their schools/districts offered staff, teachers and students guidance about use of AI. Schools in high-poverty areas were much less likely than schools in low-poverty areas to report having provided guidance.^[26] Why does it matter? Because precisely the same technology is a scaffold in one setting and a shortcut in another, depending on how it is introduced and circumscribed. Source checks, scale and appropriate reliance on the tool tend to be learned in privileged settings. AI shortcuts tend to be learned in deprived settings. While an apparent individual development deficit,

this isolation may seem to discount the impact; in truth, the cognitive stunting has the capacity to become the mechanism through which institutional inequality becomes cognitive inequality. Over time.

4 When AI Strengthens Learning Rather Than Replaces It

Any serious discussion needs to grapple with the most potent counterargument that generative AI and similar technologies can increase access, boost productivity, customize explanations and partly compensate for some longstanding weaknesses of the educational system. That case has more than rhetorical strength. UNESCO projects a looming global deficit of 44 million teachers by 2030,^[27] a shortage so large that some degree of learning support technology is no longer a luxury but a necessity. In settings with limited, delayed or uneven lesson delivery, an LLM capable of identifying the specific difficulty in a question could still provide support that would otherwise be unavailable, and could do so at scale. This counterargument should not be dismissed; to do so would be to misunderstand both the educational crisis and the allure of AI.

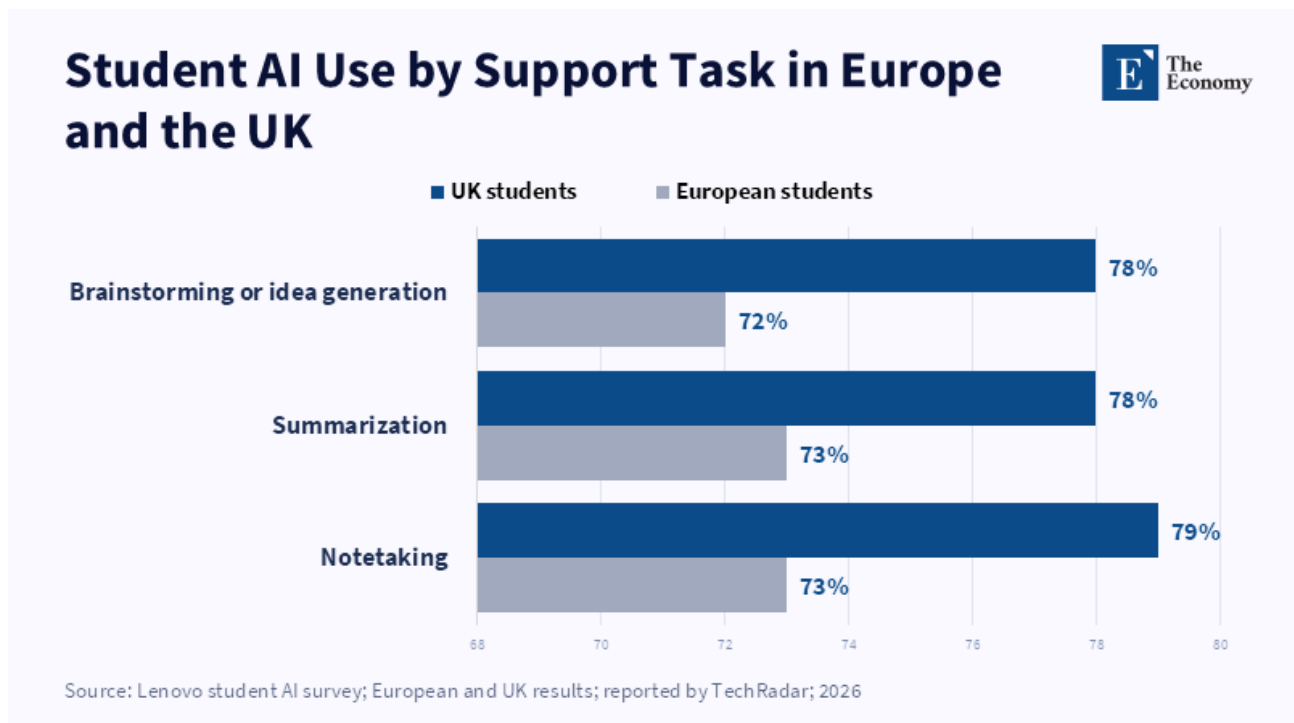


Figure 7: Students most often report using AI for support tasks such as notetaking, summarization and idea generation, not only direct answer production.

The empirical evidence of benefit is now strong enough to matter. That a 2025 randomized-control trial in Scientific Reports showed a custom AI tutor built on research-informed pedagogy achieved larger learning gains in less time^[28] than an in-class active learning condition in a real college physics context; that a 2026 paper in npj Artificial Intelligence showed that targeted fine-tuning specific to grade-levels improved the accuracy of model explanations relative to prompt-based approaches by 35.64%,^[29] without sacrificing accuracy; that, in context of German primary school math classrooms, a 2025 study with 114 students showed that context-personalized AI-generated materials and problems boosted intrinsic motivation, interest and learning outcomes;^[30] and that, in context of large language model use by undergraduate students, that a 2026 meta-study of 66 experiments on ChatGPT use showed mixed results with a generally positive effect on learning outcomes^[31] (although the

field is still nascent). The overall lesson: under some circumstances, AI can facilitate achievement.

But the same evidence also makes transparent the constraints of the optimistic interpretation. Those successful applications are almost never cases of open-ended, general-purpose, unsupervised interaction. The tutor in the physics RCT was custom-engineered for that task. The grade-specific model explicitly taxed readability and developmental level. The elementary school personalization project incorporated individual task design into the classroom activity. Those Stanford coding trials likewise demonstrate that benefits accrue when the learner employs the model as a tutor and backfire when the model does the work for them. The positive results thus do not negate the concern with cognitive offloading. They delineate the circumstances in which offloading is agentic rather than substitutive. AI works as an educational additive as long as it functions as a framework for learner effort. It weakens the system as soon as it becomes a frictionless answer-bot that disconnects performance from understanding.

This is also the point at which the familiar analogy with calculators both seems fair and yet feels short-sighted. Calculators did not undermine the study of mathematics because school systems eventually learned to sequence their use: number sense and concept understanding first, the use of the calculator later, selectively. Search engines did not undermine historical argumentation, because good teaching continued to require source comparison, argumentation and evidentiary judgment. Generative AI can be incorporated with the same ethos, but only to the extent that education systems can determine which skills will always need to be prioritized. UNESCO's advice on AI literacy^[32] is clear that regulation should be right for children, appropriate for the age, family and culture; protective of privacy; pedagogically valid, not technocratic. That's a wise position. The best argument for AI will come when schools do not abdicate to the tool, but determine beforehand which abilities the tool may supplement and which must still be personally provided.

5 Redesigning Education for an AI-Mediated Learning Environment

The policy conclusion at the heart of this paper follows directly from this analysis. The worst thing is not that AI has suddenly handed students the ability to be perpetually lazy, but that the education system still depends on assignment formats, epistemic assumptions and compliance procedures that were already inadequate and are now public knowledge. High schools and colleges have too often conflated artifact creation and idea cultivation in the past; machine-generated responses make that error transparent. Once a machine can generate a coherent, professional-quality article or program at near-zero additional expense, an education system that prizes the best work without regard for process lapses becomes untenable. The impulse to outlaw AI is therefore institutionally weak. A school can block a website for a term but it cannot restore the epistemic conditions of January 2020.

A reinvented system must use the abundance of machine assistance to redefine the purpose of education so that it yields, must, in other words, answer the question: what should education now be for? The answer cannot be romantic insistence that performance unaided by tools is always more worth engaging in and that the task of the classroom is always to forge the individual's capability in all domains, per Dewey. Today's job

and citizenship often call for dexterous and imaginative application of machines. But there can be no wholesale subcontracting of the how. If, as the IMF says, advanced-economy cognitive work is a bellwether of widespread AI-dense occupations and if the OECD cautions that generative AI's urban-rural and regional gaps may intensify, then the added value of durable human judgment must, paradoxically, go up. Reading, writing, quantitative reasoning, knowledge of one's specialist discipline, the ability to compare information sources efficiently and to test for source veracity, metacognitive judgment and a capacity for model reasoning will all be more, not less, essential for students today precisely because they will have such immensely powerful yet ever-present tools at hand. The truly scarce expertise in such a world will not be text production; it will be the ability to articulate questions with exquisite precision, weigh up diverse evidence, calibrate what is known, uncertain and contested and call on alternative assumptions, causal models and arguments of others when making explicit, publicly accountable choices under uncertainty. This is a pedagogy, not a superfluous detail.

Curriculum design must therefore become more disciplined about sequencing. For novices, an underpowered effort still counts. The most compelling findings indicate that having the trainee work on the problem, produce a draft interpretation, or articulate an explanation before using AI assistance preserves the generative power on which learning depends. Then the AI can be used in a referencing, arguing, or providing a diagnosis mode, for side-by-side comparison, critique, alternative hypotheses, guided prompts, or reduction of complexity. That is the choice between replacement and extension of cognition by the AI. Harvard research on the K-12 implementation of coding practices bears out this warning and applies in the German K-12 study on adaptation guidance, demonstrating how more reflective chatbot usage of learners can be trained without creating more overload. To put it concretely, such ordering implies a pedagogy of phased encounters: retrieval or creation first, then guided engagement and finally referential human reflection. Such ordering requires more rigor on the part of instructors than permissive use, but it yields more dignity and substance.

The work itself should also change. A system for learning that can survive AI will depend less on the generic prompt-response template that a general public model can answer in a flash and more on tasks that demand situated judgment. This does not require gimmickry. The tasks should foreground work that depends on local data, group conversation, lived experience, lab or field notes, oral explanation, collaborative decision-making, visualization and iterative drafts. Multimodal, context-specific assignments are not immune to AI, but they make one-off outsourcing less practical and more transparent and they are more true to the kind of work that we generally imagine "real" intellectual work to be and which, in fact, is never only producing an original five-paragraph paper from thin air. The goal is not to trap students, but to ensure assessment measures the kind of thinking that institutions profess to value.

Assessment redesign, therefore, appears inevitable. The literature is arguably pointing in this very direction. A 2024 scoping review in the *International Journal of Educational Technology in Higher Education* concluded that the use of generative AI calls for assessment redesign focused on self-directed learning, competence and integrity, professional development and multi-methods assessment amid,^[33] not through, detectors. A 2026 ICT-in-tertiary-education framework similarly argues: Assessment design in tertiary ICT education must be reconsidered, not patched. The implications are hugely practical. Assertions with high-stakes implications about individual learning need to lean on in-lesson tasks, oral exams, e-portfolios with version histories, su-

pervised problem-solving and reflective explanations of when and why an AI tool was used. Process evidence is vital. So is openness. Institutions ought to depart from a simple: AI is cheating versus AI is not cheating language to common-sense labels: AI-disallowed, AI-limited, AI-allowed and disclosed, AI-expected.^[34] This is not bureaucratic. It provides the lowest common denominator for consistent standards.

For administrators, the last three years have not been an epidemic of wild openness. HEPT's 2026 survey (where, lest we forget, almost all undergrad AI use is already mainstream) found only 36% of students feel their institution actually encourages AI usage and only 38% report tools and resources available. RAND reached similar conclusions (also derived from the Great Lakes States) that guidance focused on heavily policed, evenly slow implementation, particularly in high-poverty schools. That delivers the worst of all possible worlds: ubiquitous, yet unchecked, usage and no intelligible scheme for what that should look like. Institutions should therefore pursue governance structures that seem both simple and proportionate. These should (at the very least) incorporate clear, institution-wide course-level policies (as newness may require), standard (re)disclosure language, measures for safeguarding privacy and security, redress processes against misconduct, distinctions for assessing developmental learning versus summative assessment, discipline-specific training and abandon plans that rely primarily on AI detectors. Even outside of the classroom, Nature Human Behavior has found that AI detectors tend to punish academic writing that is flattish and very aware of stylistic conventions.^[35] The right institutional response is not to patch weak assessment design with unreliable language-detection tools but to improve task design and evidence substantiation.

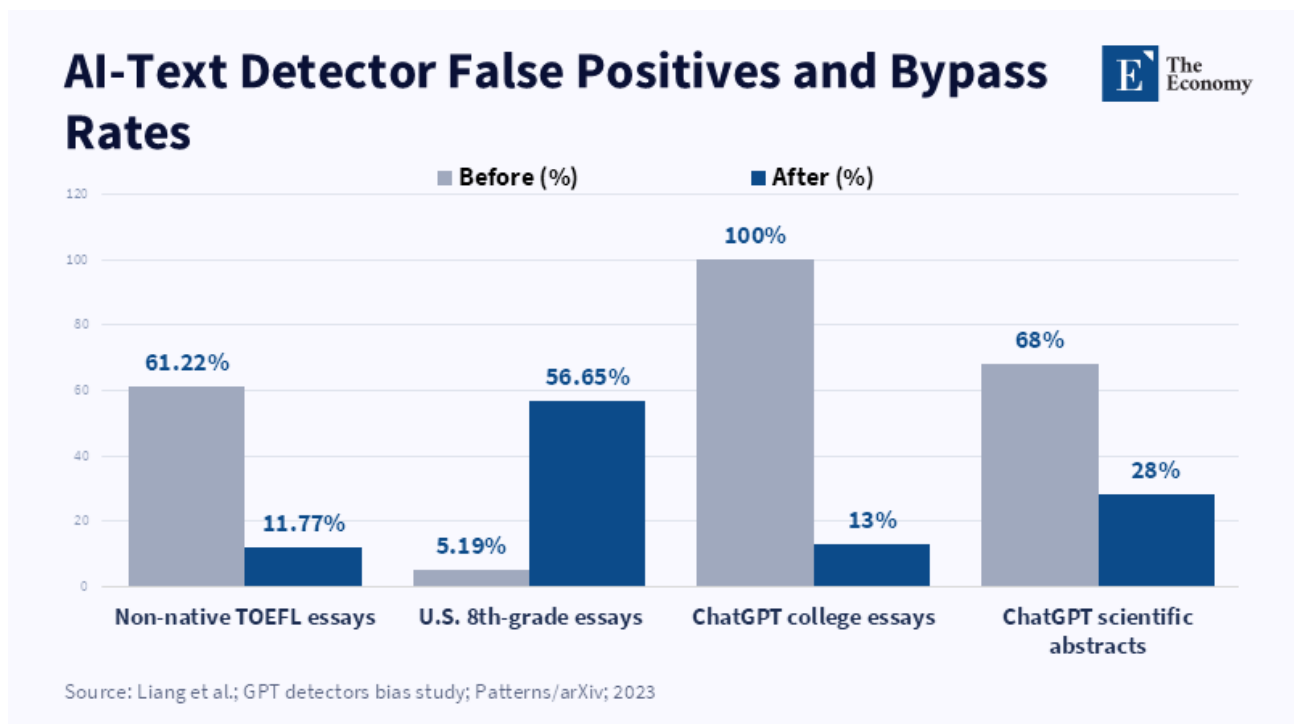


Figure 8: Detector-based enforcement is fragile because both human writing and AI-generated writing can be misclassified under small changes in style or wording.

A still broader response is required for policymakers. Education ministries and public systems cannot respond to the AI challenge as a technical or narrowly ed-tech question, because AI is already caught up with teacher vacancies, labor-market shifts and equality of opportunity. UNESCO's headcount estimate makes it clear why

support for AI is tempting and what lies behind the risk of technology substituting for investment in teachers. AI can leverage the relative scarcity of expertise; it cannot replace the human relational, motivational and ethical tasks of instruction. Public policy should therefore continue with both capacity-building and safeguarding agendas—investing in capacity, championing shared quality resources, supporting evidence-based design and procurement and developing local evaluation capability—at the same time as protecting students, establishing age-specific protections and requiring evidence on all assessed systems. UNESCO’s draft principles seem to point in this direction. Resilience should be the goal: systems that know where to leverage AI and where to limit it.

6 Education Systems, Equity and Cognitive Resilience

Equity requires special attention. The OECD’s regional analysis cautions that generative AI exposure is far greater in urban than in rural labor markets and risks widening divisions. RAND’s new evidence from schools suggests the same might be happening in education systems. If the richer schools and regions get higher quality training, safer tools and stronger governance, then AI will support existing human capital stocks—“complementing” beyond. If the poorer schools merely absorb unmanaged exposure, then AI will more often substitute for a human capital foundation that was not built in the first place—that’s the real stratification challenge. It’s not so much that some children have access to AI and others don’t: it’s that some know how to govern AI and others do not; in other words, the distinction is between those who will learn to govern AI, as opposed to those who will simply be governed by AI, which is why public provisioning is vital. If governments leave AI entirely to the market, the technology will likely widen existing disparities.

A third form of system renewal is about scholarship itself. Our intuitions that weakly generated AI material should be banished from academic life seem supported by the evidence. It is not only student work. A thirteen-fold jump in hallucinated citations (2014-26) has been identified in Nature’s 2026 reporting and major audits of scholarly repositories; tens of thousands of invalid references were identified in 2025 publication and repository archives. The use of LLM outtakes in scientific writing has been examined by Nature Human Behavior. These issues matter for education in that learners are formed by knowledge institutions as well as classrooms. When the knowledge archive begins to fill with synthetic yet untrustworthy content, verification becomes more critical. Policy should thus head for the enforcement of citation verification, provenance-aware workflows and process accountability in research training.^[36] The answer for low quality via AI is not anti-technology rhetoric. It is tighter quality controls where the knowledge archive is produced.

What emerges from all this is a more disciplined, not a more nihilistic, conclusion. AI is not a license to write less, reason less and work less carefully. There is much more reason to finally recover what we’ve always supposed to have been doing all along. Writing should be more process-led, more evidence-based and more defensible. Assessment should embody deep understanding, not simply a facility for text production. The curriculum should teach first-order content and second-order judgment about tools. Heavy-handed administrators should build governance, not improvisation. Policymakers should fund public capacity rather than leaving AI integration to private platforms alone. Education isn’t saving its dignity by performing old traditions with more gusto.

It's saving its dignity by becoming more honest, at the intellectual level, about what learning is, what we've imagined the tools do and what human capacities do and don't stay human.

7 Conclusion - Institutional Renewal, Not Panic or Surrender

In its strongest form, the "cognitive stunting" argument is not that generative AI is uniquely pernicious or that children were thriving until it arrived. It is that an already weakened education system, whose gaps are now easier to exploit, is now being challenged with a technology that can emulate the evident product of learning without the considerable unseen effort and goes beyond search engines into synthesis, formulation and advocacy. If anything, that is why it cannot be dismissed as simply Google with a better interface. The threat to learning is early and unbounded outsourcing of hard cognitive work, especially for novice writers, atop systems which continue to conflate completed products with knowledge.

The policy answer should therefore be neither panic nor complacency. AI can broaden access, tailor explanations and bolster systems under strain, particularly where shortages are most severe. Yet, the data also show that the payoffs are greatest where AI is pedagogically limited and least when it substitutes for cognitive effort. Education systems should therefore redefine curriculum, pedagogy, assessment and governance along those lines. What is needed is neither prohibition nor capitulation. It is the renewal of institutions that safeguard durable mental capabilities, teachers learners to command powerful tools and guarantee that AI supports, rather than weakens, the formation of human judgment, a matter on which future equality and social resilience will depend.

References

- [1] Favero, L., Pérez-Ortiz, J.-A., Käser, T. and Oliver, N. (2025) 'Do AI tutors empower or enslave learners? Toward a critical use of AI in education', *arXiv preprint*.
- [2, 10] Risko, E.F. and Gilbert, S.J. (2016) 'Cognitive offloading', *Trends in Cognitive Sciences*, 20(9), pp. 676–688.
- [3] Freeman, J. (2025) *Student Generative AI Survey 2025*. Oxford: Higher Education Policy Institute and Kortext.
- [4] Higher Education Policy Institute and Kortext (2026) *Student Generative AI Survey 2026*. Oxford: Higher Education Policy Institute and Kortext.
- [5] Pew Research Center (2025) *What Teens Say They Do and Don't Use ChatGPT For*. Washington, DC: Pew Research Center.
- [6] World Bank, UNESCO, UNICEF, USAID, FCDO and Bill & Melinda Gates Foundation (2022) *The State of Global Learning Poverty: 2022 Update*. Washington, DC: World Bank.
- [7] OECD (2023) *PISA 2022 Results, Volume I: The State of Learning and Equity in Education*. Paris: OECD Publishing.

- [8] International Monetary Fund (2024) *Gen-AI: Artificial Intelligence and the Future of Work*. Washington, DC: International Monetary Fund.
- [9] OECD (2024) *Job Creation and Local Economic Development 2024: The Geography of Generative AI*. Paris: OECD Publishing.
- [11] Sparrow, B., Liu, J. and Wegner, D.M. (2011) ‘Google effects on memory: cognitive consequences of having information at our fingertips’, *Science*, 333(6043), pp. 776–778.
- [12] Hutchins, E. (1995) *Cognition in the Wild*. Cambridge, MA: MIT Press.
- [13, 14, 15] Pew Research Center (2025) *Do People Click on Links in Google AI Summaries?* Washington, DC: Pew Research Center.
- [16, 36] Zhao, Z., Wang, Y., Stuart, T., De Vaan, M., Ginsparg, P. and Yin, Y. (2026) ‘LLM hallucinations in the wild: large-scale evidence from non-existent citations’, *arXiv preprint*.
- [17] OpenAI (2023) ‘New AI classifier for indicating AI-written text’, *OpenAI*.
- [18] Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. and Wilson, N. (2025) ‘The impact of generative AI on critical thinking: self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers’, *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*.
- [19] Bjork, R.A. and Bjork, E.L. (2011) ‘Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning’, in Gernsbacher, M.A., Pew, R.W., Hough, L.M. and Pomerantz, J.R. (eds) *Psychology and the Real World: Essays Illustrating Fundamental Contributions to Society*. New York: Worth Publishers.
- [20, 32] UNESCO (2023) *Guidance for Generative AI in Education and Research*. Paris: UNESCO.
- [21] Bischof, L., Schön, E.-M., Rauschenberger, M. and Neumann, M. (2026) ‘Knowing the rules is not enough: student regulatory awareness and use of GenAI in higher education’, *arXiv preprint*.
- [22] Lehmann, M., Cornelius, P.B. and Sting, F.J. (2024) ‘AI meets the classroom: when do large language models harm learning?’, *arXiv preprint*.
- [23] Koriat, A. and Bjork, R.A. (2005) ‘Illusions of competence in monitoring one’s knowledge during study’, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(2), pp. 187–194.
- [24, 25] Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X.-H., Beresnitzky, A.V., Braunstein, I. and Maes, P. (2025) ‘Your brain on ChatGPT: accumulation of cognitive debt when using an AI assistant for essay writing task’, *arXiv preprint*.
- [26] Gallup and Walton Family Foundation (2026) *Teachers and AI Guidance in K–12 Education*. Washington, DC: Gallup and Walton Family Foundation.

- [27] UNESCO, International Task Force on Teachers for Education 2030 and Fundación SM (2024) *Global Report on Teachers: Addressing Teacher Shortages and Transforming the Profession*. Paris: UNESCO.
- [28] Dai, X., Wen, Z., Jiang, J., Liu, H. and Zhang, Y. (2025) ‘How students use AI feedback matters: experimental evidence on physics achievement and autonomy’, *arXiv preprint*.
- [29] Oh, J., Whang, S.E., Evans, J. and Wang, J. (2026) ‘Classroom AI: large language models as grade-specific teachers’, *arXiv preprint*.
- [30] Harvey, E., Koenecke, A. and Kizilcec, R.F. (2025) ‘Don’t forget the teachers: towards an educator-centred understanding of harms from large language models in education’, *arXiv preprint*.
- [31] Huang, J., Wang, R.R., Liu, J.-H., Xia, B., Huang, Y., Sun, R. and Xue, J.M. (2025) ‘A meta-analysis of LLM effects on students across qualification, socialisation, and subjectification’, *arXiv preprint*.
- [33] Perkins, M., Furze, L., Roe, J. and MacVaugh, J. (2024) ‘The AI Assessment Scale in action: a pilot implementation of GenAI supported assessment’, *Journal of University Teaching and Learning Practice*.
- [34] Perkins, M., Roe, J. and Furze, L. (2024) ‘The AI Assessment Scale revisited: a framework for educational assessment’, *arXiv preprint*.
- [35] Liang, W., Yuksekgonul, M., Mao, Y., Wu, E. and Zou, J. (2023) ‘GPT detectors are biased against non-native English writers’, *Patterns*, 4(7).