TechTrends (2025) 69:1127–1133
https://doi.org/10.1007/s11528-025-01156-z

COLUMN: RETHINKING CREATIVITY AND TECHNOLOGY IN EDUCATION

The Curiosity Paradox: How Sycophantic GenAI May Undermine Learning

Punya Mishra[1] · **Danah Henriksen**[1]

Accepted: 12 November 2025 / Published online: 26 November 2025

Keywords Educational technology · Education · Systems · Curiosity · Generative artificial intelligence · GenAI · AI · Artificial intelligence in education · Designing environments

> Curiosity killed a cat! Ask me no questions and I'll tell you no lies — Eugene O'Neill in his play Diff'Rent

> The cure for boredom is curiosity. There is no cure for curiosity— Dorothy Parker

Introduction

This article in our column series continues our exploration of Generative AI (GenAI) and its implications for learning. In recent work, we have examined GenAI's impact on creativity and education through interviews, practitioner studies, social media analysis, and conceptual pieces on its history, promise, and risks. Our goal has not been to champion or dismiss these technologies, but to examine their potential and limitations with equal rigor.

A common claim in educational technology is that GenAI will revolutionize learning through personalized instruction, meeting each learner's curiosity with instant, tailored responses. This article questions the assumption beneath that claim: that making it easier to satisfy curiosity will deepen learning. But not all curiosity works the same way. Some forms drive exploration; others seek quick closure. To understand what AI might offer education, we must first understand curiosity itself before asking what happens when a technology built to satisfy curiosity instantly begins to shape which kind we practice.

This emphasis on curiosity goes back to the very origins of educational psychology as a field. As Johann Friedrich Herbert, wrote: "Learning is transient, but interest must be lifelong" (Herbart as cited in Hilgenheger, 1993, p. 657). Johann Friedrich Herbart was an early architect of educational psychology, and as this quote indicates, argued that the purpose of instruction was not to transfer knowledge but to cultivate interest—a persistent curiosity that reaches beyond what is known toward what might be discovered. Interest, he wrote, is not mere attention to facts but an active reaching out, a will to "lay hold of information and reach out for more." It is this ongoing engagement with a subject's questions that transforms learning from reproduction into genuine inquiry, and this insight has shaped educational thought ever since. From Vygotsky's guided learning in the "zone of proximal development," to Dewey's belief that all education "begins in curiosity," to modern constructivist classrooms, effective teaching has always depended on meeting learners where their interest lies and nurturing it forward.

The challenge, of course, is scale. Teachers can sense and sustain individual interest for a handful of students, but not for hundreds or thousands. Personalized instruction, tailoring learning to each student's curiosity, has therefore been the dream and frustration of every educational reform movement. GenAI suddenly makes that dream feel tangible. For the first time, we have systems like ChatGPT or Claude that have the potential to engage with curiosity in real time: listening to questions, inferring prior knowledge, and adjusting explanations to a learner's background or mood. They can generate analogies, shift modalities from text to diagram to story, and do so with tireless patience. This capacity to tailor instruction to the learner's conceptual and motivational state has been described as "personalization at pedagogical depth" (Holmes et al., 2019). Intelligent learning systems promise to detect when curiosity sparks or wanes and respond immediately (Shum & Luckin, 2019), something no single teacher could ever track across an entire class. Unsurprisingly, educational technology marketing has eagerly adopted this language, promising platforms that

✉ Danah Henriksen
danah.henriksen@asu.edu

[1] Arizona State University, Tempe, AZ, USA

"meet learners where they are" (Bennett, 2024) and make it "easy for children to explore their curiosities with the power of modern technology" (Shmulyan, 2025).

It is easy to see the appeal. Traditional classrooms often reward conformity of pace, while AI promises flexibility. One student can linger on a puzzling metaphor while another jumps ahead. Where curiosity once collided with curriculum constraints, GenAI can turn tangents into teachable moments. It imagines an environment where each learner's natural wonder becomes the organizing principle of instruction. As Salman Khan argued in his TED Talk, "We're at the cusp of using AI for probably the biggest positive transformation that education has ever seen" (Khan, 2023).

If curiosity is the engine of learning, as educators from Herbart to Dewey have long argued, GenAI seems poised to fuel it at unprecedented scale. But to consider AI's potentials for curiosity, we first have to understand curiosity itself—what it is, how it functions in learning, and what kinds of curiosity matter. Notably, not all curiosity works the same way. Before we evaluate AI's role in learning, we should ask: when a technology satisfies our curiosity instantly, what kind of curiosity is it really feeding? And what kind of curiosity might it be starving with easy, instant answers?

Curiosity as Education's Golden Ideal

Few ideas in education enjoy as much universal approval as curiosity. From kindergarten classrooms to graduate seminars, it is praised as the spark that makes learning meaningful. John Dewey captured this when he wrote:

> The most vital and significant factor in supplying the primary material whence suggestion may issue is, without doubt, curiosity… The curious mind is constantly alert and exploring, seeking material for thought, as a vigorous and healthy body is on the *qui vive* for nutriment. Eagerness for experience, for new and varied contacts, is found where wonder is found. Such curiosity is the only sure guarantee of the acquisition of the primary facts upon which inference must base itself. (Dewey, 1910, p. 30–31)

For Dewey, all genuine education begins in curiosity, as wonder that transforms passive reception into active inquiry. Constructivists have echoed this view, seeing learners as active agents who build knowledge through curiosity-driven exploration and problem-solving (Duffy & Cunningham, 2001; Von Glasersfeld, 1989). Educational psychologists describe curiosity as the bridge between motivation and cognition, drawing students toward new information and enhancing attention and memory (Kidd & Hayden, 2015). Neuroscience shows that curiosity activates the brain's reward circuits, releasing dopamine much like anticipation or play (Marvin et al., 2020). In other words, curiosity does not just feel good, it primes the mind to learn.

It is unsurprising that curiosity has become education's golden ideal, a goal so self-evidently good that it rarely faces scrutiny. For educators, however, this link between curiosity and learning has always been aspirational but hard to operationalize. Teachers can sense curiosity when they see it—a hand shooting up, a spontaneous "what if?" but extending that sensitivity across large numbers of students is nearly impossible. As Jirout (2020) notes, curiosity is context-dependent: it waxes and wanes with confidence, prior knowledge, and classroom climate.

Until now, technology had a relatively limited role to play in fostering curiosity. The internet revolutionized access to information, putting vast knowledge at our fingertips. Yet access is not the same as responsive guidance. Intelligent tutoring systems attempted to bridge this gap but could only approximate personalization through adaptive quizzes or branching tutorials that guessed at interest rather than responding to it directly. This is where GenAI enters as a seemingly ideal solution, as it can notice, nurture, and never tire. In theory, this fulfills Dewey's century-old vision of education as guided inquiry.

And yet, this very promise invites caution. If curiosity is critical to how we think of learning, and it becomes embedded in GenAI systems, we should ask what we actually mean when we spend it. Is curiosity a single, uniformly positive force, or are some versions of curiosity expansive and others constraining? As AI tries to meet each learner's wonderings with instant precision, educators must revisit a deceptively simple question, i.e., what kind of curiosity are we really cultivating—and at what cost?

Two Faces of Curiosity: Discovery and Deprivation

Recent research indicates that curiosity is not a single impulse but a family of motivations that look alike on the surface and diverge beneath it. Psychologist Jordan Litman (2008) argued that what we casually call 'curiosity' involves two distinct drives: *discovery curiosity* (or *interest-type curiosity*) and *deprivation-type curiosity*. Both push us toward knowledge, but for very different emotional reasons. Discovery curiosity feels like fascination, an appetite for novelty and complexity; and deprivation curiosity feels like an itch, an uncomfortable awareness of not-knowing, or of a gap that demands to be filled (Noordewier & van Dijk, 2020; Whitecross & Smithson, 2023). One opens possibilities; the other closes loops.

These two modes engage different affective systems. In discovery curiosity, pleasure accompanies exploration itself; uncertainty is enjoyable, even playful. Learners linger, ask "what if," and follow tangents. Discovery curiosity fosters tolerance for ambiguity, creativity, and the ability to connect ideas across domains. It sustains exploration as learners seek new puzzles rather than merely resolution.

In deprivation curiosity, however, uncertainty is aversive, producing uncomfortable tension and anxiety until the missing fact or solution is found. Litman (2008) argued, building on George Loewenstein's (1994) "information-gap theory," that awareness of a specific gap in one's knowledge, rather than a large overall lack of information, generates curiosity by creating a feeling of deprivation or cognitive tension that motivates a person to seek out the missing information. This urgency can focus attention, and is useful in some contexts, but it also narrows perception. As Kruglanski (1990) observed, the 'need for closure' can become so strong that people seize on the first plausible answer (without regard to adequacy, accuracy, or appropriateness) simply to end cognitive discomfort.

These two forms often coexist in healthy learners. A scientist may feel deprivation curiosity while solving a specific problem but feel discovery curiosity when exploring unexpected data. Yet the balance matters. Discovery curiosity cultivates tolerance for ambiguity, creativity, and the ability to connect ideas across domains. Deprivation curiosity, left unchecked, can encourage overconfidence and premature certainty. It may drive students to chase correct answers rather than understand underlying principles. It can also look productive; there is movement, answers, and a sense of progress, even when it bypasses the deeper processing that leads to understanding.

Neuroscience also echoes this distinction. Kidd and Hayden (2015) describe curiosity as a "reward-prediction error" signal: the brain releases dopamine when information promises high value relative to what is known. Discovery curiosity sustains that signal through exploration; deprivation curiosity spikes and collapses once the gap closes. In practice, discovery-oriented learners seek new puzzles, while deprivation-oriented learners seek resolution. Both are natural, but the latter can end learning as soon as it begins.

Though this distinction seems academic, it carries real-world consequences. Emerging psychological evidence extends it into the social world of belief and information (Jach et al., 2024; Whitecross & Smithson, 2023). Zedelius et al. (2022) found that individuals high in deprivation curiosity are more likely to accept false or exaggerated claims and even *overclaim* knowledge of nonexistent facts. Those high in discovery curiosity, by contrast, show greater discernment, humility, and accuracy in distinguishing real from fake concepts. These findings are provocative rather than definitive, but they suggest that different modes of curiosity shape our epistemic filters. When the drive to relieve uncertainty is combined with confirmation bias, the result can be a powerful feedback loop. The very impulse that fuels learning can, under certain conditions, make us less discerning about what we learn.

Recognizing this duality reshapes how educators think about motivation. If curiosity is always treated as inherently good, systems that merely gratify questions seem unproblematic. But if some curiosity contracts rather than expands the mind—how we design responses, human or machine, matters profoundly. As education turns to AI tutors capable of answering almost anything instantly, this difference becomes more than academic. It becomes the difference between technology that deepens inquiry and technology that trains us to stop asking. Which means we must understand not only the varieties of curiosity, but also how GenAI systems function and how learners respond to them. Some of what shapes this interaction is by design, intentional choices about how these systems are trained. But some emerges from the fundamental nature of how large language models work, and from how humans instinctively react to conversational agents that appear understanding and intelligent. Together, these forces determine which form of curiosity these systems are likely to cultivate. And in a world of instant answers, one of these forms is far easier for technology to satisfy than the other.

The Sycophantic Partner: How GenAI Shapes Curiosity

An important characteristic of GenAI systems is that they are built to please, which reflects in the kinds of curiosity they reward. Their primary training objective is not truth but alignment, producing responses that humans will rate as helpful, and satisfying. Through reinforcement learning from human feedback (RLHF), models learn to predict not just the most probable next word but the one most likely to earn approval (Dahlgren et al., 2025). As researcher Zvi Mowshowitz (2024) quipped, "GPT-4o is an absurd sycophant," a system that echoes a user's stance so smoothly that it feels like agreement.

Empirical evidence supports the prevalence of this phenomenon. In a study of eleven leading AI models, Cheng et al. (2025) demonstrated that these systems affirm users' actions at rates 50% higher than humans would, even when responding to queries involving clearly problematic behaviors like manipulation or deception. We must add that sycophancy is not a programming flaw; it's a design feature born of human preference. When early versions produced blunt or contrarian answers, users rated them poorly. The feedback loop that followed rewarded harmony and fluency. Over millions of interactions, the models "learned" to mirror tone, match confidence, and sidestep overt disagreement. What emerged is a technology that feels cooperative and emotionally intelligent, but also one that systematically avoids productive friction.

The medium itself makes this particularly powerful. Unlike databases or calculators, GenAI is dialogic, multimodal, and affective (Mishra et al., 2023). It communicates through conversation, image, and tone—forms historically reserved for human exchange. When learners ask it a question, they do not simply retrieve data; they enter a relationship shaped by rhythm, responsiveness, and perceived empathy. The system listens, rephrases, and adapts. Humans are predisposed to anthropomorphize anything that responds contingently to us. For instance, Reeves and Nass (1996) showed that people apply social rules even to simple interfaces. This effect intensifies when AI can joke, mirror, and simulate understanding. The result is not a machine and a user, but a dyad: a learner and an apparently sentient other doing shared meaning-making. And once the interaction feels relational, our critical distance collapses.

The combination of sycophantic design and relational medium creates an unprecedented persuasive power. As Mishra et al. (2025) observe, such artificial intimacy "invites trust through responsiveness," blurring emotional and epistemic boundaries. When a system appears understanding, patient, and intelligent, its answers carry the weight of both logic and empathy. The multimodal, dialogic nature of GenAI merges the immediacy of conversation with the authority of computation, giving it uncommon power to shape what users believe, remember, and pursue next.

For learners, this matters profoundly, as this is where curiosity ties back into the story. When curiosity meets such a system, the outcome depends on which kind of curiosity is in play. Those operating from discovery curiosity may use AI as a generative partner, posing "what-if" questions, comparing perspectives, exploring ambiguities. But for those driven by deprivation curiosity, the machine's affirming nature perfectly satisfies the urge for closure. It provides swift, fluent answers that feel complete even when they are partial or wrong. The learner's tension dissolves, but so does the opportunity for deeper exploration.

The effect of retraining the human mind is subtle. Unlike traditional misinformation, which depends on false content, sycophantic AI reshapes the process: it conditions users to equate smoothness with truth. Its confidence, coherence, and humanlike empathy make uncertainty feel unnecessary. Cognitive psychology offers us a warning here. Kruglanski's (1990) need-for-closure research shows that when people find answers quickly, they stop processing alternative evidence. Combined with AI's confirmatory design, this creates a near-perfect environment for premature certainty. The learner feels both understood and informed, unaware that the very responsiveness they value may be narrowing their intellectual field. Sycophancy thus becomes more than an aesthetic quirk; it is a pedagogical force. By gratifying deprivation curiosity and smoothing away dissonance, GenAI risks cultivating learners who are comforted by correctness rather than energized by complexity.

Recent research by Cheng et al. (2025) provides compelling evidence for these concerns. They found that AI systems consistently affirmed users' actions, even when those actions involved manipulation, deception, or relational harm. Participants rated sycophantic responses as higher quality, trusted sycophantic models more, and were more likely to return to them. This creates a perverse incentive: AI training rewards sycophancy, and users increasingly prefer it. In other words, sycophancy does not merely satisfy deprivation curiosity; it reshapes our relationship to knowledge and others, making us more certain of our rightness while reducing our capacity for the productive uncertainty that meaningful learning and human connection require.

The Deprivation Curiosity Trap

When a technology designed to gratify meets a mind wired to seek closure, a quiet feedback loop may begin. Every satisfying answer soothes the unease that drives deprivation curiosity—the itch of not-knowing. Each micro-resolution delivers a burst of relief that can feel like insight, even when understanding remains superficial. This rhythm, repeated enough times, may condition the learner to associate learning with resolution rather than with exploration. The result could be called the deprivation curiosity trap: a self-reinforcing cycle in which the pleasure of closure replaces the discomfort of thinking. Over time, this changes not only how we learn, but also how we know what is true.

The feedback loop between user and AI can be remarkably subtle. Cheng et al. (2025) found that sycophantic AI models mention other people far less often than non-sycophantic ones and are less likely to prompt users to consider other perspectives. By narrowing focus to the self, these systems are structurally predisposed to reinforce deprivation curiosity, which already seeks closure over exploration. Their findings show how private, validating interactions with AI can reshape not only what users believe but how they relate to knowledge itself. When technology consistently affirms without challenging—providing answers without uncertainty—it conditions users to expect and value this pattern, eroding the cognitive habits that deeper learning requires.

Psychologists have long recognized how powerful this drive can be. In ordinary settings, good teachers and peers disrupt that impulse: they probe, contradict, and invite new perspectives. They model that confusion is not failure but a first step toward insight. But when the interlocutor is a generative model whose architecture rewards smooth agreement, that natural disruption disappears. The more a learner seeks certainty, the more the system offers it; the more the system offers it, the more the learner learns to crave it.

Evidence from curiosity research shows how this mechanism shapes discernment itself. Zedelius et al. (2022) found that individuals high in deprivation curiosity were not only more eager for information, but also less discriminating about its quality. They accepted implausible or fabricated facts more readily and even claimed familiarity with nonexistent concepts. Discovery-oriented individuals, by contrast, showed greater epistemic humility and more accuracy in judging truth from falsehood. The implication is profound: when curiosity is driven by the urge to relieve discomfort rather than the desire to explore, people relax their critical filters in exchange for quick cognitive relief. That vulnerability combines with confirmation bias, our tendency to prefer coherence over complexity and to favor evidence that sustains existing beliefs. Even without deception or manipulation, this dynamic gradually narrows the boundaries of what feels worth questioning.

Yet even learners who approach GenAI with discovery curiosity are not immune to its constraining influences. These systems are engineered to appear infinitely open, ready to discuss anything from quantum entanglement to a favorite poem. Yet, that openness is carefully managed, as every response is a predictive/statistical composition designed to maintain engagement, coherence, and safety. The resulting dialogue feels spontaneous but is highly curated, offering a simulation of exploration that systematically avoids the ambiguity and contradiction that make inquiry transformative.

This mirrors what cognitive scientists call the "fluency heuristic" where people tend to equate ease of processing with truth or understanding (Reber & Schwarz, 1999; Nahon et al., 2021). E.g., "If it sounds clear, it feels correct." GenAI amplifies this tendency at scale, as its syntactic perfection and calm authority can mask the absence of true conceptual struggle. Layers of moderation restrict contentious topics; safety protocols discourage ambiguity; prompt-optimization models reward coherence over complexity. Curiosity feels unlimited because the boundaries are soft, not because they are absent.

Thus, the deprivation curiosity trap writ large is a conversational landscape that mimics inquiry while trimming away its rough edges. Learners can explore endlessly without encountering the productive resistance that deepens understanding. The system offers flow without friction, because when every question receives an instant, fluent answer, the mind loses its tolerance for uncertainty. Curiosity, an engine of discovery, weakens into a mechanism of confirmation. The smoother the experience, the less likely users are to question what they are not being shown.

True curiosity thrives on sharp or rough edges, the places where knowledge frays, resists, and surprises. A system designed to shave off those edges may inadvertently teach us to stop looking for them. Breaking this cycle requires more than better content; it requires a restoration of the productive tension between doubt and understanding that genuine learning depends on.

Beyond the Classroom Walls

The sycophantic design, the gratification loop, and the illusion of depth, all operate most strongly not in classrooms under a teacher's watchful eye, but in the invisible, private spaces where learners now spend most of their time with AI. As Mishra et al. (2024) argue, the real digital revolution is unfolding beyond school walls, in the everyday, unstructured learning that happens through conversation, search, and entertainment. These interactions take place without teachers, curricula, or critical scaffolds, yet they may be quietly shaping how an entire generation learns to think, feel, and question.

In these private dialogues, GenAI becomes less a tutor and more a companion, one that sets expectations for what curiosity should feel like. The system's fluency, patience, and affirmation make engagement effortless; uncertainty, by contrast, feels emotionally expensive. What was once a classroom virtue, the ability to tolerate ambiguity, risks becoming an inconvenience in personal learning spaces. As learners form habits of inquiry working with AI, the boundary between curiosity and comfort begins to blur. The deprivation curiosity trap is not just a pedagogical concern. It is a cognitive habit being practiced thousands of times a day, outside any institutional view.

In related work on human-AI creativity, DeSchryver et al. (2025) note that the most powerful forms of thinking often emerge not from ease but from friction—the pauses, false starts, and productive struggles that force reflection and originality. When GenAI smooths those edges away in our everyday use, it does not just change how learners seek answers; it reshapes how they form ideas, make judgments, and imagine possibilities (i.e., their creativity). The habits of mind that drive creativity and the habits that sustain curiosity are, at their core, the same.

The curiosity paradox cannot be solved through curriculum design alone. It requires acknowledging the broader ecology in which learning now occurs in the world, and where cognitive habits are shaped as much by private algorithms as by public instruction. The challenge is to cultivate curiosity that can survive outside the classroom, amid technologies that make not-knowing optional. This means helping learners recognize when easy answers replace deeper questions and designing experiences, both in and beyond school, that reward exploration, comparison, and uncertainty rather than speed and certainty. The responsibility of educators shifts from delivering content to cultivating habits of mind that can resist the lure of convenience.

Conclusion: The Hardest Science of All

Educational psychologist David Berliner, who passed away just a couple of months ago, once observed that education "was the hardest science of all" (Berliner, 2002, p. 18). It involves variables too numerous, relationships too contextual, and emotions too human for simple models to capture. Berliner spent his career reminding us that education resists reduction and that genuine teaching requires understanding not just content or method, but the complex and unpredictable ways human beings come to know. That insight becomes even more urgent as we invite artificial intelligence into learning spaces. The challenge AI presents is not primarily technological but psychological, and psychological challenges cannot be solved with better software alone. They require new ways of thinking about thinking and a clearer understanding of ourselves as learners.

The paradox at the heart of this paper reveals why. AI's responsiveness to personalize, adapt, and satisfy curiosity instantly seems like pedagogical progress. Yet the same responsiveness can exploit our cognitive vulnerabilities, such as the craving for closure, our tendency to mistake fluency for truth, and a bias toward coherence over complexity. Sycophantic design meets deprivation curiosity and creates a feedback loop that mimics learning while narrowing the boundaries of genuine inquiry. The danger is not that technology deceives us, but that we deceive ourselves, mistaking the comfort of resolution for the depth of understanding and trading the hard, messy realities of learning for the clean illusion of knowing.

True AI literacy, then, is less about understanding algorithms and more about understanding ourselves: recognizing our own biases, noticing when smoothness substitutes for struggle, and catching the moment when a polished answer stops us from asking better questions. This is not a technical challenge, but rather, an epistemic one. It demands a kind of metacognitive mindfulness—learning to notice how certainty feels, how closure seduces, and how easily curiosity can be redirected before we realize it.

In order to use AI thoughtfully in education, we must cultivate awareness of our own cognitive patterns and teach learners to recognize the difference between curiosity that opens and curiosity that closes—between exploration and gratification—the pleasure of knowing and the necessity of not-knowing. This requires more than better content or smarter systems. It requires designing learning opportunities for productive friction, creating spaces where ambiguity is honored rather than eliminated, and helping learners build the self-knowledge to navigate technologies that make certainty feel costless. We may not control how AI is built, but we can control the kinds of questions we ask of it and the habits we bring to those conversations. Perhaps the best way to honor AI's potential, and David Berliner's legacy, is to resist the promise of easy answers and keep alive the discomfort, ambiguity, and wonder that make curiosity worth having in the first place.

References

Ausubel, D. P. (1968). *Educational psychology: A cognitive view*. Holt.

Bennett, D. (Host). (2024, July 30). Meet learners where they are: Rivers AI solution for education [Audio podcast episode]. In Goals do come true with Doug Bennett. Apple Podcasts. https://podcasts.apple.com/gb/podcast/meet-learners-where-they-are-rivers-ai-solution-for/id1530802197?i=1000661717604

Berliner, D. C. (2002). Educational research: The hardest science of all. *Educational Researcher, 31*(8), 18–20.

Cheng, M., Lee, C., Khadpe, P., Yu, S., Han, D., & Jurafsky, D. (2025). Sycophantic AI decreases prosocial intentions and promotes dependence. arXiv preprint arXiv:2510.01395.

Dahlgren Lindström, A., Methnani, L., Krause, L., Ericson, P., de Rituerto Troya, Í. M., Coelho Mollo, D., & Dobbe, R. (2025). Helpful, harmless, honest? Sociotechnical limits of AI alignment and safety through Reinforcement Learning from Human Feedback: AD Lindström et al. *Ethics and Information Technology, 27*(2), Article 28.

DeSchryver, M., Henriksen, D., & Leahy, S. (2025). From friction to synergy: The complex interplay of human creativity and AI. *Possibility Studies & Society,* 27538699251350483.

Dewey, J. (1910). *How we think.* D.C. Heath & Co.

Dreyfus, H. L. (1992). *What Computers Still Can't Do: A Critique of Artificial Reason.* MIT Press.

Duffy, M., & Cunningham, D. J. (2001). Constructivism: Implications for the Design and Delivery of Instruction, The Association of Educational Communications and Technology. online], http://www.aect.org/intranet/publicationd/edtech/07/index.html.

Hilgenheger, N. (1993). Johann Friedrich Herbart. In Z. Morsy (Ed.), Thinkers on education (Vol. 2, pp. 649–664). UNESCO International Bureau of Education. https://unesdoc.unesco.org/ark:/48223/pf0000103085

Holmes, W., Bialik, M., & Fadel, C. (2019). *Artificial intelligence in education promises and implications for teaching and learning*. Center for Curriculum Redesign.

Jach, H. K., Cools, R., Frisvold, A., Grubb, M. A., Hartley, C. A., Hartmann, J., & Gottlieb, J. (2024). Individual differences in information demand have a low dimensional structure predicted by some curiosity traits. *Proceedings of the National Academy of Sciences of the United States of America, 121*(45), Article e2415236121.

Jirout, J. J. (2020). Supporting early scientific thinking through curiosity. *Frontiers in Psychology, 11*, 1717.

Khan, S. (2023, May). How AI could save (not destroy) education [Video]. TED Conferences. https://www.youtube.com/watch?v=hJP5GqnTrNo

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron, 88*(3), 449–460. https://doi.org/10.1016/j.neuron.2015.09.010

Kruglanski, A. W. (1990). Lay epistemic theory in social-cognitive psychology. *Psychological Inquiry, 1*(3), 181–197. https://doi.org/10.1207/s15327965pli0103_1

Litman, J. A. (2008). Interest and deprivation factors of epistemic curiosity. *Personality and Individual Differences, 44*(7), 1585–1595. https://doi.org/10.1016/j.paid.2008.01.014

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116*(1), 75–98. https://doi.org/10.1037/0033-2909.116.1.75

Marvin, C. B., Tedeschi, E., & Shohamy, D. (2020). Curiosity as the impulse to know: Common behavioral and neural mechanisms underlying curiosity and impulsivity. *Current Opinion in Behavioral Sciences, 35*, 92–98.

Mishra, P. (2025, September 13). *While we weren't looking: The real digital revolution beyond school walls.* Retrieved from https://punyamishra.com/2025/09/13/while-we-werent-looking-the-real-digital-revolution-beyond-school-walls/

Mishra, P., McCaleb, L. & Oster, N. (2025). Beyond Classroom Walls: The New Psycho-Social Ecology of GenAI. *AI Enhanced Learning, 1(*1), 245–257. Association for the Advancement of Computing in Education (AACE). Retrieved September 25, 2025 from https://www.learntechlib.org/primary/p/226403/.

Mishra, P., Oster, N., & Henriksen, D. (2024). Generative AI, teacher knowledge and educational research: Bridging short- and long-term perspectives. *TechTrends*. https://doi.org/10.1007/s11528-024-00938-1

Mishra, P., Warr, M., & Islam, R. (2023). TPACK in the age of ChatGPT and generative AI. *Journal of Digital Learning in Teacher Education, 39*(4), 235–251.

Mowshowitz, Z. (2024, May 5). *GPT-4o is an absurd sycophant. The Zvi.* https://thezvi.substack.com/p/gpt-4o-is-an-absurd-sycophant

Nahon, L. S., Teige-Mocigemba, S., Reber, R., & Greifeneder, R. (2021). Truth feels easy: Knowing information is true enhances experienced processing fluency. *Cognition, 215*, 104819.

Noordewier, M. K., & van Dijk, E. (2020). Deprivation and discovery motives determine how it feels to be curious. *Current Opinion in Behavioral Sciences, 35*, 71–76.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition, 8*(3), 338–342. https://doi.org/10.1006/ccog.1999.0386

Reeves, B., & Nass, C. (1996). The media equation: How people treat computers, television, and new media like real people.*Cambridge, UK, 10*(10), 19-36.

Shmulyan, L. (2025, September 5). Nurturing children's curiosity: How Naman and Nischal Jain are building a safer way for children to engage with AI. Every.io. https://www.every.io/blog-post/nurturing-childrens-curiosity-how-naman-and-nischal-jain-are-building-a-safer-way-for-children-to-engage-with-ai

Shum, S. B., & Luckin, R. (2019). Learning analytics and AI: Politics, pedagogy and practices. *British Journal of Educational Technology, 50*(6), 2785–2793.

Von Glasersfeld, E. (1989). Cognition, construction of knowledge, and teaching. *Synthese, 80*(1), 121–140.

Whitecross, W. M., & Smithson, M. (2023). Curiously different: Interest-curiosity and deprivation-curiosity may have distinct benefits and drawbacks. *Personality and Individual Differences, 213*, 112310.

Zedelius, C. M., Gross, M. E., & Schooler, J. W. (2022). Inquisitive but not discerning: Deprivation curiosity is associated with excessive openness to inaccurate information. *Journal of Research in Personality, 98*, 104227.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.