

**Artificial Intelligence and Education:
A Neurocognitive Perspective**

Peter C. Seelman

Johns Hopkins University

AS.190.423 Planetary Geopolitics

Professor Daniel Deudney

19 May 2025

I. Close to Home

I invite you to walk with me through the library at Johns Hopkins University. We enter through the glass doors of Brody Cafe, and the aroma of freshly brewed coffee wafts through the air. To your left, tables and chairs, and students filling every nook and cranny. Each with a laptop—sometimes an iPad too—in front of them. Movement in the corner of your eye catches your attention. Your eyes dart over to a girl's screen and notice text rapidly appearing. She hastily writes down the equations now in front of her. Splashes of color—*of personality*—interrupt the flurry of text and mathematics. She pauses briefly and chuckles at the corny math joke that just materialized, then gets back to work.

You now notice that the guy next to her has the same dark interface on his screen too. He furiously types in a question, hits enter, and a wall of text punctuated by code blocks is built before him. After a short prompt, an image is generated on another girl's screen three tables down—it's of a character on a tropical beach, somewhere far, far away from the library. It's got a pretty impressive resemblance to her, except that it is in the unmistakable lush, ostensibly hand-drawn, and vibrant style of Studio Ghibli.

It's finals season at Hopkins, and everyone is cramming for their exams. The thing on everyone's screens? AI programs: OpenAI's ChatGPT, Google's Gemini, Anthropic's Claude, and others. This scene is representative of a phenomenon that is happening in libraries, classrooms, offices, and homes all around the world.

A flurry of questions enters your mind: Is there any historical technological precedent? What does this mean for education? What does this software do to our brains? What's happening cognitively? What are the dangers? The benefits? How can we best utilize them?

And maybe, just maybe, *is this humanity's last exam?*

II. How Did We Get Here? A Brief History of Cognitive Tools

Circa 370 BCE, Plato, in his dialogue Phaedrus, recounted a conversation between the Egyptian god Theuth, the inventor of letters, and King Thamus. Theuth presented his invention with great promise:

"This invention, O king," said Theuth, "will make the Egyptians wiser and will improve their memories; for it is an elixir of memory and wisdom that I have discovered."

But Thamus was skeptical, warning of unintended consequences:

"For this invention will produce forgetfulness in the minds of those who learn to use it, because they will not practice their memory... You have invented an elixir not of memory, but of reminding; and you offer your pupils the appearance of wisdom, not true wisdom..." (Plato, trans. 1917, Phaedrus, 274e-275b).

A. Unpacking Thamus's Warning: A Perennial Concern

Thamus's ancient fears—the potential erosion of innate human memory and the seductive allure of superficial understanding—resonate through the ages, echoing a perennial anxiety that often accompanies the introduction of transformative cognitive tools. Yet, history reveals a nuanced reality. Writing, despite Thamus's prescient concerns about the nature of recall versus true understanding, did not merely lead to forgetfulness; it became the bedrock of civilization. It enabled the recording of complex laws, the development of sophisticated philosophical arguments, the preservation and transmission of intricate narratives across generations, and the very organization of increasingly complex societies. It wasn't just an offloading of memory; it was a fundamental transformation in the scaffolding of human thought and societal structure.

B. The Lineage of Cognitive Tools: From Papyrus to Pixels

Writing was arguably the first great cognitive tool to externalize a core mental function, but it was far from the last. The trajectory of human innovation is marked by a succession of such tools, each extending our intellectual reach and reshaping our cognitive landscape. The printing press in the 15th century democratized access to written knowledge, catalyzing intellectual revolutions and societal change. Later, mechanical calculators automated arithmetic, freeing human minds for more complex mathematical reasoning. The digital age brought personal computers, the internet, and powerful search engines like Google. These tools placed vast information at our fingertips, fundamentally altering how we seek, process, and retrieve knowledge. Each innovation, while offering immense benefits, also subtly (or not so subtly) reshaped cognitive habits and raised new questions about human reliance on external aids.

C. Enter AI: A New Chapter, or a Different Book?

Artificial Intelligence, particularly the recent surge in powerful generative models, marks the latest—and arguably most disruptive—chapter in this ongoing narrative of cognitive tool development. Its perceived utility is evident in its rapid adoption across education and myriad industries. For instance, by early 2025, ChatGPT alone reported reaching 400 million weekly active users (Backlinko, 2025; Exploding Topics, 2025a). In educational settings, the uptake has been remarkably swift; one 2025 survey indicated that 92% of higher education students were using AI in some form, a significant increase from 66% in the previous year (Higher Education Policy Institute, 2025). Similarly, by May 2024, 46% of U.S. teachers reported regularly using AI chatbots, representing an annualized growth rate of 250% from Fall 2023 (NinjaED, 2025).

Beyond academia, a significant 78% of global companies reported integrating AI into their business operations by 2025 (Exploding Topics, 2025b).

This rapid proliferation has been accompanied by intense discussion about the trajectory of AI development itself. While progress in recent years has been remarkable, some experts and industry observers have raised concerns about a potential "plateau" in AI capabilities, particularly for the Large Language Models (LLMs) that have driven much of the recent excitement (Ironclad, 2025; Hypertxt, 2024). Arguments for such a slowdown include diminishing returns from scaling up model size and training data (VKTR.com, 2025; Quantilus, 2024), the challenge of exhausting high-quality public training datasets (Advon Commerce, 2025), perceived architectural limitations in achieving robust reasoning or avoiding factual errors (Hypertxt, 2024), and the substantial costs of frontier model development (VKTR.com, 2025). Some analyses suggest that while LLMs have become adept at many tasks, scaling reasoning capabilities presents unique complexities not easily solved by more data or parameters alone (Zhang et al., 2025).

However, this perspective is by no means universal. Counterarguments highlight that performance on new and more demanding benchmarks continues to show significant improvement (Stanford HAI, 2025). Moreover, the field is actively pivoting from a primary focus on sheer scale towards innovation in model efficiency, the strategic use of specialized or proprietary data, the development of novel architectures, and the creation of more sophisticated agentic systems capable of complex multi-step reasoning and tool use (Ironclad, 2025; Advon Commerce, 2025; Hypertxt, 2024). Active research into areas such as advanced algorithmic discovery, new architectures, and the exploration of highly specialized applications in nuanced domains, such as AI-assisted therapeutic tools are indicative of this continued dynamism

(Novikov et al., 2025; DeepMind, 2025, Sakana AI, 2024; Seelman and Kroumov, 2025). These avenues of research, which we will explore in more detail when considering the qualitative distinctions of AI's capabilities, suggest that rather than a definitive or singular plateau, the field may be diversifying and entering new phases of innovation.

Ultimately, the debate over whether AI capabilities will continue their recent exponential trajectory or encounter significant near-term ceilings, while intellectually stimulating, may be somewhat moot given how disruptive AI already is. Nevertheless, understanding this debate underscores the dynamic nature of AI development; it highlights that even if progress were to slow, the current capabilities are already significant enough to warrant the educational considerations explored in this paper. The technology's sheer prevalence, advanced capabilities (despite limitations), and rapid embedding across domains demand intense consideration of whether we face a societal and cognitive shift as seismic as that brought by writing.

Indeed, what is not open to debate is the present reality: even if AI development were to encounter an unforeseen barrier tomorrow, its existing capabilities and rapidly expanding usage are already exerting a considerable global impact, the full dimensions of which are still unfolding. While the debate on AI's developmental trajectory continues, its current capabilities already mark a significant departure from prior cognitive tools. This qualitative difference in capability is crucial to understand.

D. The Crucial Distinction: A New Threshold in Cognitive Automation and Interaction

While the historical lineage of cognitive tools reveals a consistent pattern of externalizing mental loads, and the rapid adoption rates of contemporary Artificial Intelligence are one indicator of its perceived utility, the more profound distinction lies not just in *usage scale* but in

the *nature of the tasks* AI can now address. We are witnessing a fundamental shift from tools that primarily aided in information storage, retrieval, and calculation, to systems demonstrating emergent capacities for complex problem-solving, novel content generation, and even sophisticated reasoning within specialized domains.

This transformation is vividly illustrated and significantly amplified by AI's burgeoning multimodal capability—the capacity to understand, process, and generate information across various formats like text, images, audio, and even video. This allows for richer, more intuitive interactions and moves AI far beyond a text-in, text-out paradigm. For instance, tools like Google's NotebookLM, powered by models such as Gemini, can ingest a wide array of source materials—PDFs, websites, and even YouTube video transcripts or audio files—and then generate outputs like an "Audio Overview," a podcast-style discussion of the provided sources (Google, 2025a; Google, 2024). While limitations certainly exist, especially with highly abstract or specialized content (as personal experience with subjects like advanced physics can confirm), the potential to alter how users engage with and synthesize diverse information is clear.

This cross-modal prowess is coupled with State-Of-The-Art (SOTA) performance in visual reasoning. OpenAI's o3 and o4-mini models, for example, demonstrate an ability to "think with images in their chain-of-thought," not merely processing them as data inputs (OpenAI, 2025a; OpenAI, 2025b), achieving high accuracy on challenging multimodal benchmarks like MMMU (evaluating college-level visual problem-solving), MathVista (visual math reasoning), and CharXiv-Reasoning (scientific figure reasoning) (OpenAI, 2025a). Such capabilities underscore that AI is tackling tasks requiring deep, specialized knowledge and complex reasoning. This push into advanced problem-solving is further evidenced by achievements like Google DeepMind's AI systems (AlphaProof and AlphaGeometry 2) reaching silver-medal

performance in the International Mathematical Olympiad (Unite.AI, 2024; AI Toolhouse Blog, 2024), and findings that GenAI can partly solve assignments at a Master's level, albeit with variable quality (Mozelius et al., 2024).

Indeed, the capacity for AI to contribute to, and even drive, discovery in highly specialized fields is becoming increasingly apparent. DeepMind's AlphaEvolve, for instance, has demonstrated an ability to go beyond optimizing existing code to discover entirely novel and more efficient algorithms. It achieved a significant breakthrough by finding an algorithm for 4x4 complex-valued matrix multiplication using only 48 multiplications, improving upon a long-standing benchmark, and has also yielded practical optimizations in real-world systems, such as enhancing data center scheduling and refining tensor processing unit (TPU) circuit designs at Google (Novikov et al., 2025; DeepMind, 2025). Concurrently, architectural innovations aim to foster more robust and adaptable AI cognition. Sakana AI's development of Cognitive Thinking Machines (CTMs), through techniques like evolutionary model merging, seeks to combine multiple specialized AI models to create more versatile systems capable of enhanced mathematical reasoning and nuanced language understanding, aspiring to a more human-like adaptability and problem-solving prowess (Uchida et al., 2024; Sakana AI, 2024).

Therefore, AI is no longer just an "elixir of reminding," as Thamus feared of writing, but is increasingly capable of acting as an "elixir of thinking," or at least, a powerful simulator of complex cognitive processes. Its ability to generate novel content (from algorithms to conversational dialogues), construct arguments, and solve intricate problems across multiple modalities signifies a departure from merely augmenting existing human skills. It points towards a technology that could potentially reshape the very nature of intellectual labor and skill development, even extending into nuanced areas like emotional well-being, with startups such as

HumorHealer aiming to complement traditional therapies (Seelman and Kroumov, 2025). This deeper impact on how we think, learn, create, and even relate, despite ongoing concerns about accuracy and potential biases (Alliant International University Center for Teaching Excellence, 2025), is what marks AI as a distinct and revolutionary step in the evolution of cognitive tools.

E. Outsourcing Reason Itself?

This unprecedented capability to simulate and automate complex cognitive tasks brings us to a more unsettling parallel than even Thamus's warning about writing. The Roman philosopher Seneca, in his *Epistulae Morales*, described a wealthy contemporary, Calvisius Sabinus. Sabinus, possessing a poor memory but desiring a reputation for intellect, purchased educated slaves—one to memorize Homer, another Hesiod, and so on—believing that "what his slaves knew, he knew too" (Seneca, trans. 1917, *Epistulae Morales*, 27.5-7). Sabinus outsourced memory, a feat Thamus might have recognized with concern. But as we witness AI tools not just recalling information but constructing arguments, solving complex problems, and generating creative text across multiple modalities, the provocation sharpens: *If Seneca's contemporary outsourced memory to his slaves, are we, with AI, beginning to outsource reason itself?* This fundamental question about the nature of the cognitive tasks we are offloading and its implications for our own intellectual development leads directly to a deeper examination of cognitive offloading in the age of AI.

III: Cognitive Offloading: Are We Still Thinking?

"My worry for my kids, my worry for society is creating technologies that make it incredibly alluring to automate the part of creation that is most difficult, most laborious, and most likely to lead to genuine insight and the sharpening of your own mental acuity... We better fucking hope that the AIs can make autonomous innovations because we are going to stop being able to."

– Ezra Klein on the Lex Fridman Podcast, 2025

A. Defining Cognitive Offloading with Established Technologies

Cognitive Offloading—the use of external aids to alter the information processing requirements of a task so as to reduce cognitive demands—has only recently become the target of systemic investigation when applied to novel digital technologies (Risko & Gilbert, 2016; Gerlich, 2025). Historically, the literature primarily concerned itself with various aspects of memory, and indeed, the definition is sometimes framed predominantly in such terms (Gerlich, 2025, p. 2). This historical focus is understandable. Modern digital tools like internet search engines, and more recently sophisticated AI, present revolutionary alternatives for offloading a broader array of cognitive functions. These range from information storage, retrieval, and complex spatial navigation to critical thinking—the central concern of this paper.

A growing body of evidence suggests that these modern technological tools often reduce cognitive burden by automating routine tasks, allowing users to delegate memory, attention, and even some decision-making processes. Perhaps one of the most widely discussed examples is the "Google Effect," a term describing our societal reliance on the internet as an externalized memory store. Seminal research by Sparrow et al. (2011) demonstrated that when individuals expect to have future access to information online, they exhibit lower rates of recall for the

information itself but enhanced recall for *where* to access it. This suggests a fundamental shift in how we manage information, prioritizing knowledge of access routes over internal storage—a modern-day, democratized analogue to Seneca's Calvisius Sabinus outsourcing his memory to enslaved individuals. Subsequent research has further explored how this constant access might lead to a blurring of boundaries between internal and external knowledge, potentially impacting how deeply information is processed and fostering a sense of knowing that is dependent on external retrieval cues (e.g., Ward, 2013; Fisher et al., 2015). While this offloading can be efficient, it raises questions about the depth of understanding and the cognitive skills developed when a quick search often supplants more effortful internal recall and integration.

Navigation, a fundamental cognitive function involving planning routes, reading maps, identifying landmarks, and maintaining direction, relies on multiple cognitive and perceptual processes (Yavuz et al., 2023; Newcombe, 2023; Spiers et al., 2023; Wolbers & Hegarty, 2010). Faulty navigation can result in distress and, in severe cases, life-threatening situations (Coughlan et al., 2018). Furthermore, deficits in wayfinding can be early indicators of conditions like Alzheimer's Disease or Traumatic Brain Injury (Coughlan et al., 2018; Seton et al., 2023).

The proliferation of GPS-based navigation systems offers a clear case of cognitive offloading. Instead of actively engaging the cognitive abilities required for traditional wayfinding, relying on GPS often means delegating these processes to external devices. Numerous studies have suggested that frequent or passive use of GPS systems may be detrimental to human navigation performance, potentially leading to poorer spatial memory and a diminished ability to form accurate cognitive maps of an environment (Yavuz et al., 2023). It has been proposed that this occurs because individuals do not sufficiently exercise and develop their innate navigational skills when consistently relying on GPS (He & Hegarty, 2020).

However, the narrative is not entirely one-sided. Interestingly, not all studies have shown that reliance on GPS is uniformly detrimental. Some research suggests that how one interacts with GPS devices is a crucial factor. Actively engaging with the device—for instance, by paying attention to the map display, anticipating turns, and relating the GPS information to the surrounding environment—may lead to better navigation performance and spatial learning compared to passively following turn-by-turn directions (Yavuz et al., 2023). GPS devices can also enrich the experience of navigating new environments under certain conditions (Yavuz et al., 2023).

Furthermore, engagement with other forms of technology, such as action video games, has been associated with *improved* spatial navigation performance and visuospatial competencies (Yavuz et al., 2023; Milani et al., 2019; Alghadari et al., 2023). This suggests that the cognitive impact of technology is not predetermined by the technology itself but is significantly shaped by the nature of the user's interaction and the specific cognitive demands of the technological task.

This distinction—that *how one uses technology* like GPS, and the specific demands of that technology (e.g., active problem-solving in video games versus potentially passive following in some GPS use cases), have dramatic implications for neurocognitive effects—is critical. It highlights that the relationship between technology use and cognitive skills is complex and mediated by user engagement. This principle of active versus passive engagement and its differential cognitive outcomes becomes even more salient and multifaceted when we turn to the latest generation of cognitive tools: Artificial Intelligence, which offers a far broader spectrum of cognitive tasks that can be potentially offloaded.

B. The Spectrum of Offloaded Cognition: What Are We Handing Over?

The capabilities of contemporary Artificial Intelligence, as explored earlier, extend far beyond simple information lookup or calculation. This proliferation into daily and professional life (Gerlich, 2025; Renner & Chaban, 2024) has, perhaps inevitably, opened up an unprecedented spectrum of cognitive tasks that can now be outsourced to these sophisticated digital tools. Inherent in this proliferation is the opportunity to offload cognitive tasks onto AI, thus necessitating the usage of a more expansive definition and analysis of how current capabilities fit into that.

The most intuitive form of offloading remains *memory delegation*. AI systems, ranging from ubiquitous search engines like Google to AI-enhanced note-taking applications (such as Notion AI) and conversational LLMs, readily serve as external repositories for declarative or episodic information. This reduces our reliance on internal biological memory processes for encoding, storage, and retrieval (cf. Risko & Gilbert, 2016; Sparrow et al., 2011), mirroring the "Google Effect" on a grander scale. Beyond simple recall, AI also takes on significant *computational burdens*. Complex mathematical calculations, quantitative data analysis, statistical modeling, and symbolic manipulations, traditionally demanding focused mental effort, can now be largely delegated to major LLMs capable of generating and executing code for data analysis (e.g., using Python), solving equations, or specialized platforms like Wolfram Alpha.

Furthermore, the initial, often laborious, stages of research and learning are being transformed. AI now offers powerful means for *offloading information retrieval and initial synthesis*. Advanced LLMs, such as Perplexity AI or Gemini Advanced, alongside AI-powered academic search tools like Elicit and Scite.ai, can locate, collate, filter, summarize, and even

perform preliminary thematic analyses of vast and diverse sources, assisting the initial stages of research or learning. This capability extends to the automation of procedural tasks. The *execution of rule-based sequences, complex workflows, or the generation of procedural content*—be it computer code, step-by-step instructions, or even optimized algorithms as seen with systems like AlphaEvolve—can be efficiently handled by AI, with major LLMs playing a significant role in code generation, debugging, and drafting standardized documents.

Perhaps one of the most striking advancements is in *content generation offloading*. AI's capacity to create novel or derivative content is no longer limited to text. Multimodal LLMs (like OpenAI's GPT-4o or Google's Gemini) can produce images, audio, and video, while specialized tools generate musical compositions or custom interactive experiences. Google's NotebookLM, for example, can even create text-to-audio overviews from source materials. This creative capacity is complemented by *perceptual offloading*, where AI augments, interprets, or filters human sensory perception and the subsequent cognitive analysis of that data (Risko & Gilbert, 2016). This includes AI tools for image recognition and description (vital for accessibility), real-time language translation, medical image analysis, and AI-powered data visualization that can highlight subtle patterns indiscernible to the human eye.

The scope of offloading extends into higher-order cognitive domains, including *decision support and, increasingly, decision automation*. AI systems analyze complex situations, model scenarios, evaluate options, predict outcomes, and suggest actions across various fields—from financial advising and clinical decision support in healthcare to strategic game playing and even everyday pros/cons analysis provided by LLMs. Finally, and perhaps most intimately, we are seeing the emergence of *emotional and social interaction offloading*. Advanced conversational LLMs (such as Character.ai or Replika) and AI chatbots designed for mental well-being support

(like Woebot or HumorHealer) are being engaged for emotional support, companionship, advice on personal matters, or to simulate therapeutic interactions, thereby outsourcing aspects of human social and emotional needs.

This sweeping potential for cognitive offloading across at least eight forms means very few technologies could rival AI's capability. The internet and cellphones, already deeply intertwined in daily life, are now becoming powerful vectors for AI-driven cognitive offloading, as companies such as Apple, Google, and Samsung heavily integrate advanced AI capabilities into their ubiquitous products (Long, 2025). The global cell phone market reached 1.22 billion units in 2024, and more than 91% of Americans own a smartphone, spending an average of 4.5 hours on it and checking it 144 times daily (Bjørhovde and Chaurasia, 2025; Bazen, 2025). This usage creates a compounded, near-omnipresent opportunity for humans to delegate a cornucopia of cognitive tasks. The advancement and widespread adaptation of artificial intelligence thus presents what may be the most significant catalyst for neurocognitive alteration in modern human history. This transformative potential is not without its complexities, and thus, the following sections will critically examine its effects, especially on higher-order thinking and education.

C. The Critical Juncture: AI, Offloading, and Higher-Order Thinking

As with almost all human technology, artificial intelligence tools are not inherently good or bad. This extends to their impact on the human mind, which is multifaceted, complex, and insufficiently understood given the lack of longitudinal studies (Chen, 2025; Gerlich, 2025; Abbas et al., 2024; Singh, 2025). These tools promise enhanced efficiency, personalized experiences, and unprecedented access to information alongside all the use cases mentioned

above. However, in conjunction with these benefits, there is growing scholarly concern regarding the cognitive and social impacts of AI on human users, particularly regarding critical thinking skills (Gerlich, 2025; Abbas et al., 2024).

Critical thinking is a "multifaceted cognitive process that involves the capacity to think clearly and rationally, understand logical connections between ideas, evaluate arguments, and identify inconsistencies in reasoning" (Gerlich, 2025, p. 3). It is crucial for academic success, informed citizenship, and professional life (de Bie et al., 2015). While AI offers potential aids, such as AI-powered data visualization tools that can help users identify trends and correlations in large datasets, thereby supporting analytical thinking (Gerlich, 2025, p. 5), the risk of over-reliance looms large. As with the passive following of GPS, such over-reliance can lead to a superficial understanding of information and a diminished capacity for independent thought. This sentiment is starkly captured by a participant in a recent study: "The more I use AI, the less I feel the need to problem-solve on my own. It's like I'm losing my ability to think critically" (Gerlich, 2025). This modern concern echoes Thamus's ancient fears about writing potentially eroding innate human faculties, now quantified by recent studies such as Gerlich (2025).

Empirical evidence is beginning to quantify these concerns. Gerlich's (2025) study of 666 participants across diverse age groups and educational backgrounds found a strong positive correlation ($r = +0.72$) between AI tool use and cognitive offloading. Cognitive offloading, in turn, was demonstrated to serve as a mediating factor between AI tool usage and critical thinking ($r = -0.75$). Ultimately, the study reported a significant negative correlation ($r = -0.68$) between frequent AI tool usage and critical thinking skills. While higher education levels appeared to mitigate some of these negative effects, the overall relationship between AI tool usage and

critical thinking was found to be non-linear, with diminishing returns in critical thinking as AI reliance increased (Gerlich, 2025, p. 16).

Qualitative data from the same study, derived from 50 semi-structured interviews using Braun and Clarke's thematic analysis framework, shed further light on these dynamics. Three dominant themes emerged: AI Dependence, concerns about Skill Atrophy, and perceptions of Cognitive Enhancement. AI dependence was particularly pronounced among younger participants (17-25 age group), who described a "heavy reliance" on AI tools for tasks ranging from "simple information retrieval" to "complex decision-making processes," noting how AI assistants and enhanced search engines had become integral to their daily routines (Gerlich, 2025, p. 20).

The impact of Generative AI (GenAI) on knowledge workers' cognitive processes reveals a subtle but significant shift in effort, as observed by Singh et al. (2025). For lower-order cognitive tasks like recall and comprehension, the user's focus often shifts from active information gathering to reactive information verification. When applying information, the emphasis moves from direct problem-solving to integrating AI-generated responses. Crucially, for higher-order tasks involving analysis, synthesis, and evaluation, the human effort transitions from direct task execution to a more supervisory role of "task stewardship" (Singh et al., 2025).

This shift has nuanced implications for creativity and originality—key components of critical thinking. While GenAI can assist less creative individuals in producing more engaging outputs, as Doshi and Hauser (2024) found with writers, this often comes at the cost of "collective diversity," with AI-assisted creations tending towards homogeneity. Peschl (2024) echoes this concern, arguing that AI models risk stifling creativity by recycling existing

knowledge, thereby creating feedback loops that reinforce repetitive patterns. Furthermore, Kumar et al. (2025) found that LLMs are more effective for convergent, goal-oriented thinking tasks but can hinder divergent thinking, which requires exploration and unconventional approaches. Viewing creativity as a skill that develops through practice (Sternberg, 2006), the widespread adoption of GenAI could lead to a decline in human creative abilities as individuals reduce the exercise of their own creative faculties (Singh et al., 2025).

These findings underscore a central concern. While offloading mundane tasks to AI is beneficial, the ease of AI generating sophisticated outputs risks inadvertently offloading critical thinking itself. This includes the nuanced evaluation, deep analytical engagement, and creative problem-solving vital for genuine understanding (Gerlich, 2025; Celik et al., 2025). Given that critical thinking is paramount for effective problem-solving, informed decision-making, and knowledge acquisition, its potential erosion due to uncritical AI adoption has profound implications, particularly for education (Murawski, 2014; Gerlich, 2025). The challenge, therefore, is to navigate this new technological landscape in a way that harnesses AI's power without undermining our own cognitive vitality, a dilemma that directly impacts pedagogical approaches in the age of AI.

IV: Education in the Age of AI

Back when I was serving on my high school Student Council, the administration moved to ban virtually all phone use during the school day. There was uproar across the student body, and—as the representatives of the people—we assembled a task force to negotiate the removal of this policy.

We assembled a report replete with academic studies, a survey of the student body, and point-by-point arguments against the move. Many of my colleagues contended that phones were

not as distracting as they were chalked up to be and that their necessity outweighed any such drawbacks.

I, however, insisted that the ban was flawed precisely because phones were a constant potential distraction. If students did not learn how to work and live effectively with them now, what would happen when they went to college where there were no teachers to take them away and tell them to get back to work? My school would be failing in its college-preparatory mission! Eventually, the administration relented and we settled on a compromise: no phones during class time, but during study halls, students had a choice: get a headstart on that homework or scroll through an endless sea of TikTok videos.

This experience with mobile phones offers a pertinent analogy for the current challenges posed by AI in educational settings: outright bans are often impractical and may miss opportunities for teaching responsible use. The imperative, then as now, is to guide students in navigating these powerful tools effectively, rather than attempting futile prohibitions.

A. The Necessity for Change

Not long ago, educators believed that content knowledge was sufficient for students to succeed. It was thought that—for the most part—that the information that students *should* learn was the same information as their parents learned. The advent of modern technology, of instant access to information (accurate, biased, flat-out-wrong, or otherwise), 24/7 news cycles, hyper political polarization, rapid communication and adoption of ideas, has left this pedagogical paradigm obsolete. It became imperative that students learn how to sort through and question the sea of information at their fingertips (Murawski). This is why critical thinking should be (and is) a pillar of education, it allows students to become diligent curators, critics, and creators of new information.

Section III illustrated how using AI tools often correlates with cognitive offloading, which can lead to a decrease in critical thinking. This is particularly noticeable with students and

the younger generation writ large. But, let us take a step back for a moment. The entire point of these AI tools is that they are *tools*; they are designed to aid us in carrying out some function. Given that we have engineered this particular tool to have near limitless applicability, it is up to us to decide how, and to what end, we employ it. Many of the studies, anecdotes, and analyses above have implicitly or explicitly suggested the critical fact that a significant portion of the neurocognitive effects of AI depends entirely upon how we use this technology.

As argued in Section IIIB, the advancement and widespread adaptation of artificial intelligence seems to place humanity upon the precipice of the most significant neurocognitive alteration in modern human history. Among many other things, this necessitates another paradigm shift in education. The following section aspires to be an actionable, research-informed proposal for exactly that.

B. Standardized Education About AI

Standardized curricula like the K-12 Common Core State Standards, ACT recommendations, and Columbia's Core Curriculum aim to cultivate essential intellectual tools, critical thinking, and readiness for complex challenges (National Governors Association Center for Best Practices & Council of Chief State School Officers, 2010; ACT, 2006; Columbia University, 2025).

The prevalence, multi-domain applicability, and multi-modal nature of artificial intelligence (AI) necessitate education about appropriate usage across practically all majors or concentrations. Recognizing this, the U.S. government has signaled proactive steps, such as the April 2025 presidential action aimed at "Advancing Artificial Intelligence Education for American Youth" (The White House, 2025a). This initiative, further detailed in an accompanying

fact sheet (The White House, 2025b), underscores a commitment to fostering AI competency by promoting its integration into education from K-12 through lifelong learning. Key elements of such national strategies often include comprehensive teacher training and the formation of public-private partnerships to develop and deploy AI educational resources effectively (The White House, 2025a; The White House, 2025b). Integrating AI literacy into foundational educational frameworks, including the K-12 pipeline influenced by standards like the Common Core and higher education curricula, is thus highlighted as a national priority. This involves not only understanding AI as a technology but also developing the critical thinking and ethical reasoning skills to use AI tools responsibly and effectively—skills that initiatives like the Common Core and specialized college programs aim to foster.

While national policy and curriculum are being transformed by this technology, it is prudent for students, educators, and administrators to act now to avoid the negative consequences outlined in Section IIIC.

C. Teaching the Teachers

Effective education of the students must begin with the education of the teachers. This can be done via professional training programs, especially for pre-service teachers (Celik, 2023). It should begin by detailing available tools. Currently, chatbots, intelligent tutoring systems, dashboards, and automated assessment systems are the most prevalent AI-based technologies in K-12 education (Akgun and Greenhow, 2022; Celik et al., 2022; Chen et al., 2022). These should be coupled with the potential benefits for each technology. For instance, chatbots can be used to maintain learners' motivation (Huang et al., 2022) and extend the benefits of small class sizes and one-on-one interaction with instructors to more students. Intelligent tutoring systems allow students to be provided with learning content based on their needs (Pai et al., 2021)—thereby

reducing teachers' workload so that they may dedicate more time to instruction (Mohamed & Lamia, 2018). AI-enhanced dashboards enable teachers to monitor student knowledge construction, cognitive and emotional engagements to support learning progress (Verbert et al., 2013; Pardo et al., 2018). Lastly, automated grading systems recognize and automatically score responses from students. These systems can help teachers to assess students' writing and speaking in the exams and assignments, and can be particularly useful for studying foreign languages (Wang & Zhao, 2020).

Assessing the decisions by AI-based tools requires a different set of skills than found in common pedagogical knowledge (Buckingham, 2022). There is little understanding of how teachers interpret and evaluate AI-based decisions, and even less empirical evidence explaining how teachers' instructional skills to use AI are associated with their ethical assessment (Celik, 2023).

This gap in understanding is critical, in part because the data used in training AI-based technology has inherent biases, resulting in potentially biased outputs from AI algorithms. Consequently, AI might increase existing inequality gaps among learner subpopulations like gender, race, and socio-economic status (The Institute for Ethical AI in Education, 2021).

Indeed, recent research by Harvey et al. (2025) highlights that these are not just theoretical concerns. In their study, multiple educators reported firsthand experiences with biased AI content, including the generation of "Eurocentric" outputs and "negative stereotypes about Black children." This underscores the practical reality of biased algorithms impacting the learning environment.

Beyond biased content, Harvey et al. (2025) also found a common concern among educators that LLM-based edtech could exacerbate systemic inequality by imposing costs that schools and districts are differentially able to bear. The Covid-19 pandemic and resulting school closures have already highlighted the severe consequences of digital exclusion, where learners without sufficient device and internet access experienced significant learning losses. This detrimental impact on disadvantaged youth, which could have been mitigated, serves as a stark warning.

To prevent similar failures and address these multifaceted challenges, the ethical integration of AI tools in education is an urgent necessity (The Institute for Ethical AI in Education, 2021). Teacher training should include modules on critically evaluating AI outputs for bias, accuracy, and pedagogical suitability, moving beyond mere technological operation.

The responses of educators to biased outputs currently vary, illustrating the critical need for such targeted training. Harvey et al. (2025) found that while some teachers chose to limit LLM use as a result of encountering bias, others "expressed confidence in their abilities to mediate students' interaction with these outputs, often by creating 'teachable moments' to address bias directly." This proactive approach of turning biased outputs into "teachable moments" is a key skill that teacher training programs should aim to cultivate, ensuring educators feel equipped to guide students in critically assessing and contextualizing AI-generated information rather than simply restricting access.

A recent study indicated that teachers with greater knowledge for interacting with AI tools better understand AI's pedagogical value and can better evaluate AI decisions through technological knowledge. However, effective AI integration requires combining technological

and pedagogical knowledge (Celik, 2023). Toward this aim, a scale to measure the knowledge for instructional AI use based on the technological, pedagogical, and content knowledge (TPACK) framework was expanded to include ethical aspects. The research model suggests some practical implications for the proposed professional development courses: they should prioritize both technological understanding and pedagogical applications of AI tools to foster ethical awareness and effective integration. Teacher training should explain how AI works (e.g., natural language processing in chatbots, learning from datasets) to illuminate ethical concerns like transparency and bias. Simultaneously, it must emphasize the pedagogical benefits (e.g., personalized feedback) and demonstrate field-specific AI tools to promote practical application and technological content knowledge (Celik, 2023).

D. AI in the Classroom

Educators and policymakers should promote balanced AI integration in educational settings, ensuring that AI tools complement rather than replace cognitive tasks. Active learning strategies and critical thinking exercises can help mitigate the negative effects of cognitive offloading and support the development of essential cognitive skills (Gerlich, 2025).

Furthermore, students must develop metacognitive skills to help them become aware of when and how to use AI tools appropriately without undermining their cognitive development (Gerlich, 2025).

Consequently, the status quo, wherein many classes either have no AI policy or the technology is banned outright without effective mechanisms to enforce it, is unacceptable and detrimental to learners and teachers. Given the ubiquity and ease-of-access of the technology as outlined in Section IIC, banning AI is equivalent to banning the internet when it comes to

practicability. Results of current research may showcase how to follow the guiding principle of stimulating learners' mental engagement with rich problem-solving support while discouraging passive or superficial use (Lehmann et al., 2025).

Lehmann et al. (2025) found that students who *substitute* some of their learning activities with LLMs (e.g., by generating solutions to exercises) increase the volume of topics they can learn about but decrease their understanding of each topic. Additionally, higher confidence in generative artificial intelligence is correlated with less critical thinking (Lee et al., 2025). This is dangerous when coupled with the fact that LLMs increase perceived learning by more than can be explained by actual differences in learning (Lehmann et al.). Thus, the *worst* way students learn with AI is by using it solely for task completion. This approach means they bypass material interaction, blindly trust AI output, and then falsely believe they understand the vast content "covered." The cognitive drawbacks of this approach should be clearly taught to students, and this behavior should be explicitly deemed unethical or against policy as outlined in a class's syllabus.

The same study found that students who *complement* their learning activities with LLMs (e.g., by asking for explanations) do not necessarily increase the rate of topics covered but do increase their understanding. Also, higher self-confidence and initial knowledge combined with a healthy doubt in the AI's output is associated with increased critical thinking in students (Lee et al., 2025; Lehmann et al., 2025). This is bolstered by the current understanding of the learning process, which states that at the beginning of learning a new skill, cognitive effort and persistence are essential. Building "cognitive endurance," meaning the ability to sustain "effortful mental activity over a continuous stretch of time," is critical for deeper understanding (Singh et al., 2025). Therefore, students should be taught about the unreliability of AI output and

should have an opportunity to engage with material before practicing with artificial intelligence. Guardrails in educational AI tools—such as teacher mediation—should facilitate the AI then serving functions of providing formative feedback, ensuring the learners have the ability to exercise their own judgement and critical thinking skills when they seek AI assistance (Singh et al., 2025).

Furthermore, Lehmann et al. (2025) noticed an interesting phenomena, the ability to copy and paste when interacting with an LLM is a strong determinant of usage behavior as it substantially increases the number of times students prompt the LLM for solutions. Removing copy and paste created an effort barrier to LLM usage and thereby also limits its effect on learning. Consequently, subjects might forgo the LLM if they perceive the manual copying effort as too high. Alternatively, they might make mistakes causing incorrect LLM answers or copy solutions incorrectly, leading to similar outcomes (Lehmann et al., 2025). Students were thus more diligent and intentional with their interactions with AI, reducing the likelihood of cognitive offloading.

Students' self-control problems extend to technology use. As stated earlier, learning is inherently an effortful and demanding task. They tend to prioritize immediate rewards over the effort required for future academic benefits, demonstrating a present bias (Lehmann et al., 2025). Students may thus prefer substitution because it reduces their learning effort. The increasing capabilities of large language models and the challenge of sustaining self-control suggest that relying on voluntary guidelines may be ineffective (Gerlich, 2025). Teachers must therefore facilitate student understanding of the consequences of cognitive offloading to AI so that they may make informed choices whenever interacting with AI.

Additionally, teachers must construct assignments, curricula, and tests with a focus on what skills students will actually need in the age of AI—avoiding the overemphasized formulaic decision making that AI is naturally good at (Singh et al., 2025). This leads to a discussion about assignments.

E. Intentional Assignment and Assessment Design

Not every assignment should use AI. Given AI's capabilities (Section II), preventing its use requires creativity. Classroom-supervised, AI-free assignments are obvious but not always practicable, making analysis of AI-resistant assignments useful. To ensure authenticity in an AI-pervasive environment, educators are developing innovative assignments. NYIT's Center for Teaching and Learning (2024) offers valuable models:

- **Shifting Modalities and Emphasizing Personal Connection:** Some approaches aim to create tasks less amenable to direct AI generation. For instance, some educators suggest assignments requiring students to record themselves reading texts and then reflect on the phenomenological differences between reading and listening. Others have students watch video interviews with authors and summarize their impressions, centering the task on unique, non-textual source material and individual interpretation. Tasks that ask students to transcontextualize existing works by rewriting them in their own contemporary neighborhood or connecting them to current, specific news events demand a high degree of personalization and creative synthesis.
- **Directly Engaging with AI as a Learning Tool:** Rather than attempting to circumvent AI entirely, other strategies involve incorporating it directly into the learning process, focusing on critical engagement with AI outputs. One approach involves offering parallel

assignment versions: one AI-free (with a student declaration) and another where students use AI to generate an initial response, then critically analyze, correct, and expand upon that output using track changes; grading then focuses on the depth of this critical engagement. This is diametrically opposed to the aforementioned passive task completion as students learn to critique and actively build upon an AI output. Another strategy is to have students use AI as a 'first critic,' submitting a draft to an AI for feedback, then revising their work based on that critique, submitting both drafts and the AI's suggestions. AI can also be integrated into interactive classroom techniques like Think-Pair-Share, modifying this to "Think-Pair-Ask ChatGPT-Pair-Share," allowing students to use AI to refine their collaboratively developed ideas in real-time and practice critical assessment of AI-generated content. This last point is perhaps the most essential of all potential activities: having students develop prompts that will foster learning and actively critique generative AI outputs (Oates and Johnson, 2025; Singh, 2024)

These examples reveal a vital pedagogical evolution, where assignments are crafted not only to mitigate the pitfalls of uncritical AI use but also to strategically deploy AI in fostering critical thinking, metacognition, and profound understanding. In this framework, intentional design can transform AI into an "elixir of thinking," positioning the student as the principal architect of their learning.

While the preceding discussion outlines pedagogical frameworks, the practical application of AI by students often involves navigating a complex interplay of academic demands, available technological affordances, and individual learning ethics. To provide a concrete illustration of these common challenges, a hypothetical scenario depicting a student's engagement with AI for a demanding academic task across advanced studies is presented in

Appendix A. This narrative is intended to offer insight into the typical cognitive and situational pressures that can influence students' choices and learning experiences in the current technological landscape.

V: Conclusion

Our inquiry began in a library, one filled with students eagerly utilizing a new technology that promised to change their educational journey and the world that they live in. While seemingly novel, it is the latest development in a long lineage of cognitive tools dating back over three thousand years to an age of papyrus, kings, and gods. The adoption of each such invention promised to extend our intellectual reach and reshape our cognitive landscape. Nonetheless, the great thinkers of these times were hesitant; the tradeoffs were unclear and we worried that these technologies could change who we were—indeed, they probably did.

What *is* novel about this particular technology is its unprecedented capability for complex reasoning and generation, coupled with its rapid, widespread adaptation. This combination places us upon the precipice of what may be the most significant catalyst for neurocognitive alteration in modern human history. While cognitive offloading is an inherent consequence of such a powerful tool, as was the case with its predecessors, the crucial difference lies in the *nature* and *scale* of what can be offloaded. We have a unique opportunity to proactively manage this process. If we do so, we might be able to redirect cognitive effort towards novel forms of creativity or problem-solving not previously possible. Upon these capabilities, we could continue to build our collective human identity. If we do not, then we risk becoming reliant on these machines, and—like a modern version of Seneca's contemporary offloading his memories—giving up one of our most precious facets: reasoning.

This paper is a small contribution to the effort to realize the former option. While our understanding is still evolving, it is clear that we have the power and knowledge to enact the necessary reforms. They begin, as these things so often do, with our youth. First, we must educate them to resist the temptation of easy answers and share the value of persisting through the difficulty associated with learning. They need to know the consequences of offloading so that they can make informed decisions on when, where, and how to complement natural with artificial intelligence. We must teach our teachers how this technology works, its flaws, its strengths, ethical best practices, and how to design assignments to facilitate this transformation.

With these educational changes, the fear that we are taking humanity's last exam may indeed prove excessive. Perhaps we are only taking a crucial midterm, one for which we are, and must continue to be, well-prepared.

References

- Abbas, M., Jam, F. A., & Khan, T. I. (2024). Is it harmful or helpful? Examining the causes and consequences of generative AI usage among university students. *International Journal of Educational Technology in Higher Education*, 21, 10.
<https://doi.org/10.1186/s41239-024-00444-7>
- ACT, Inc. (2006). Assessing the college readiness in reading of eighth and ninth-grade students using ACT's EXPLORE®. ACT.
<https://www.act.org/content/dam/act/unsecured/documents/EXPLOREReading.pdf>
- Advon Commerce. (2025, January). Has generative AI hit a plateau? Exploring the next wave of AI innovation. Advon Commerce Topics.
<https://www.advoncommerce.com/topics/has-gen-ai-hit-a-plateau>
- AI Test Kitchen Blog. (2024). AlphaProof and AlphaGeometry 2: AI achieves silver-medal level in solving International Mathematical Olympiad problems.
- AI University Center for Teaching Excellence. (2025). Concerns about generative artificial intelligence.
- Akgun, S., & Greenhow, C. (2022). Artificial intelligence in education: Addressing ethical challenges in K-12 settings. *AI and Ethics*, 2(4), 431–440.
<https://doi.org/10.1007/s43681-021-00096-7>
- Alghadari, F., Alhaqbani, A., Aljammaz, M., AlMousa, R., & AlGhannam, H. (2023). Exploring the influence of video games on self-reported spatial abilities among university students. *Multimodal Technologies and Interaction*, 7(4), 90. <https://doi.org/10.3390/mti7040090>
- Al-Harrasi, S., & Al-Aufi, A. (2025, January). Critical thinking in the age of AI: A systematic review of AI's effects on higher education. *Education Quarterly Reviews*.

AutoGPT. (2025, April). 10 best AI for math problems.

<https://autogpt.net/the-best-llm-for-math-problem-solving/>

Backlinko. (2025, April). ChatGPT / OpenAI statistics: How many people use ChatGPT?

Retrieved May 19, 2025, from <https://backlinko.com/chatgpt-stats>

Barr, M. S. (2025, May). Speech by Governor Barr on artificial intelligence and the labor market.

Bazen, A. (2025, March 20). Cell phone statistics 2025. ConsumerAffairs.

https://www.consumeraffairs.com/cell_phones/cell-phone-statistics.html

Burns, L. (2025). Beyond problem-solving: The future of learning in an AI-driven world.

Creative Education, 16(4), 520–534. <https://doi.org/10.4236/ce.2025.164031>

Canalys. (2025, February). Global smartphone market soared 7% in 2024 as vendors prepare for tricky 2025. Retrieved May 19, 2025, from

<https://canalys.com/newsroom/worldwide-smartphone-market-2024>

Celik, I. (2023). Towards intelligent-TPACK: An empirical study on teachers' professional knowledge to ethically integrate artificial intelligence (AI)-based tools into education.

Computers in Human Behavior, 138, 107468. <https://doi.org/10.1016/j.chb.2022.107468>

Celik, I., Dindar, M., Muukkonen, H., & Järvelä, S. (2022). The promises and challenges of artificial intelligence for teachers: A systematic review of research. *TechTrends*, 66(4),

616–630. <https://doi.org/10.1007/s11528-022-00715-y>

Center for Teaching & Learning, New York Institute of Technology. (2024, October). ChatGPT:

The future of learning or the death of critical thinking? Retrieved May 19, 2025, from

https://site.nyit.edu/ctl/blog/chatgpt_the_future_of_learning_or_the_death_of_critical_thinking

- Chen, B. (2025). Beyond tools: Generative AI as epistemic infrastructure in education [Preprint]. arXiv. <https://arxiv.org/abs/2504.06928>
- Chen, X., Zou, D., Xie, H., Cheng, G., & Liu, C. (2022). Two decades of artificial intelligence in education: Contributors, collaborations, research topics, challenges, and future directions. *Educational Technology & Society*, 25(2), 28–47. [https://doi.org/10.30191/ETS.202204_25\(2\).0002](https://doi.org/10.30191/ETS.202204_25(2).0002)
- Chung, D. J. H., Gao, Z., Kvasiuk, Y., Li, T., Münchmeyer, M., Rudolph, M., & Tadepalli, S. C. (2025). Theoretical physics benchmark (TPBench) -A dataset and study of AI reasoning capabilities in theoretical physics [Preprint]. arXiv. <https://arxiv.org/abs/2502.15815>
- Columbia College. (n.d.). The core curriculum. Retrieved May 19, 2025, from <https://www.college.columbia.edu/core-curriculum>
- ConsumerAffairs. (2025, March). Cell phone statistics 2025. Retrieved May 19, 2025, from https://www.consumeraffairs.com/cell_phones/cell-phone-statistics.html
- Coughlan, G., Laczó, J., Hort, J., Minihane, A.-M., & Hornberger, M. (2018). Spatial navigation deficits—Overlooked cognitive marker for preclinical Alzheimer disease? *Nature Reviews Neurology*, 14, 496–506. <https://doi.org/10.1038/s41582-018-0031-x>
- de Bie, H., Wilhelm, P., & van der Meij, H. (2015). The Halpern Critical Thinking Assessment: Toward a Dutch appraisal of critical thinking. *Thinking Skills and Creativity*, 17, 33–44. <https://doi.org/10.1016/j.tsc.2015.04.003>
- Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(21). <https://doi.org/10.1126/sciadv.adn5290>
- Exams, S. M. (2025, February). Over 75% of students use AI for homework.

- Fisher, M., Goddu, M. K., & Keil, F. C. (2015). Searching for explanations: How the Internet inflates estimates of internal knowledge. *Journal of Experimental Psychology: General*, 144(3), 674–687. <https://doi.org/10.1037/xge0000070>
- Fridman, L. (Host). (2025, May). Ezra Klein and Derek Thompson: Politics, Trump, AOC, Elon & DOGE (No. 462) [Audio podcast episode]. In Lex Fridman Podcast.
- Gerlich, M. (2025). AI tools in society: Impacts on cognitive offloading and the future of critical thinking. *Societies*, 15(1), 6. <https://doi.org/10.3390/soc15010006>
- Gong, C., & Yang, Y. (2024). Google effects on memory: A meta-analytical review of the media effects of intensive Internet search behavior. *Frontiers in Public Health*, 12, 1332030. <https://doi.org/10.3389/fpubh.2024.1332030>
- Google. (2025). Google NotebookLM | Note taking & research assistant powered by AI.
- Google Blog. (2024, September). NotebookLM now lets you listen to a conversation about your sources.
- Grinschgl, S., & Neubauer, A. C. (2022). Supporting cognition with modern technology: Distributed cognition today and in an AI-enhanced future. *Frontiers in Artificial Intelligence*, 5, 908261. <https://doi.org/10.3389/frai.2022.908261>
- Halx, M. D., & Reybold, L. E. (2006). A pedagogy of force: Faculty perspectives of critical thinking capacity in undergraduate students. *The Journal of General Education*, 54(4), 293–315. <https://doi.org/10.1353/jge.2006.0009>
- Harvey, E., Koenecke, A., & Kizilcec, R. F. (2025). ‘Don’t forget the teachers’: Towards an educator-centered understanding of harms from large language models in education [Preprint]. arXiv. <https://arxiv.org/abs/2502.14592>

He, C., & Hegarty, M. (2020). Individual differences in the integration of global and local information in navigation. *Topics in Cognitive Science*, 12(3), 910–934.

<https://doi.org/10.1111/tops.12509>

Higher Education Policy Institute. (2025, February). Student generative AI survey 2025.

Holmes, L., Rios, J., Berice, B., Benson, J., Bafford, N., Parson, K., & Halloran, D. (2021).

Predictive effect of *Helicobacter pylori* in gastric carcinoma development: Systematic review and quantitative evidence synthesis. *Medicines*, 8(1), 1.

<https://doi.org/10.3390/medicines8010001>

Huang, W., Hew, K. F., & Fryer, L. K. (2022). Chatbots for language learning—Are they really useful? A systematic review of chatbot-supported language learning. *Journal of*

Computer Assisted Learning, 38(1), 237–257. <https://doi.org/10.1111/jcal.12610>

Hypertext Blog. (2024, November). AI's limits: The impending performance plateau.

Inc., I. (2025, January). What to expect in 2025: Predicting the year in AI. Ironclad Blog.

Institute for Ethical AI in Education. (2021, March). The ethical framework for AI in education.

The University of Buckingham.

Jarvee. (2025, January). 10 best AI tools for homework assignments.

Keuning, T., & van Geel, M. (2021). Differentiated teaching with adaptive learning systems and

teacher dashboards: The teacher still matters most. *IEEE Transactions on Learning*

Technologies, 14(2), 201–210. <https://doi.org/10.1109/TLT.2021.3072143>

Kumar, H., Vincentius, J., Jordan, E., & Anderson, A. (2025). Human creativity in the age of

LLMs: Randomized experiments on divergent and convergent thinking. *Proceedings of*

the 2025 CHI Conference on Human Factors in Computing Systems, 1–18.

<https://doi.org/10.1145/3706598.3714198>

LearnFast Ai. (2025). LearnFast AI: The ultimate solver for math and physics powered by advanced AI.

Lee, H.-P. (Hank), Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & 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*, 1–22.
<https://doi.org/10.1145/3706598.3713778>

Lehmann, M., Cornelius, P. B., & Sting, F. J. (2025). AI meets the classroom: When do large language models harm learning? [Preprint]. arXiv. <https://arxiv.org/abs/2409.09047>

Long, E. (2025, April 20). AI smartphone features compared—iPhone 16 vs. Samsung Galaxy S25 vs. Google Pixel 9. Tom's Guide.
<https://www.tomsguide.com/ai/ai-smartphone-features-compared-iphone-16-vs-samsung-galaxy-s25-vs-google-pixel-9>

Milani, L., Camisasca, E., Caravita, S. C. S., Gatti, E., & Tizzoni, M. (2019). Positive effects of videogame use on visuospatial competencies: The impact of visualization style in preadolescents and adolescents. *Frontiers in Psychology*, 10, 1226.
<https://doi.org/10.3389/fpsyg.2019.01226>

Mohamed, H., & Lamia, M. (2018). Implementing artificial intelligence in education for optimal results. *International Journal of Information Technology and Computer Science*, 10(2), 43–51. <https://doi.org/10.5815/ijitcs.2018.02.06>

Mozelius, P., Hellerstedt, A., Klingström, M., & Ghorbani, S. (2024). Generative AI and its impact on activities and assessment in higher education: Some recommendations from

master's students. Proceedings of the 4th International Conference on AI Research (ICAIR 2024).

Murawski, L. M. (2014). Critical thinking in the classroom...and beyond. Semantic Scholar. <https://api.semanticscholar.org/CorpusID:149417192>

National Governors Association Center for Best Practices, & Council of Chief State School Officers. (2010). Common Core State Standards for English Language Arts & Literacy in History/Social Studies, Science, and Technical Subjects. Authors. https://learning.ccsso.org/wp-content/uploads/2022/11/ELA_Standards1.pdf

Newcombe, N. S., Hegarty, M., & Uttal, D. (2023). Building a cognitive science of human variation: Individual differences in spatial navigation. *Topics in Cognitive Science*, 15(1), 6–14. <https://doi.org/10.1111/tops.12626>

NinjaED. (2025, January). AI in education statistics 2024: Teacher usage & adoption trends. Retrieved May 19, 2025, from <https://www.ninjaed.com/ai-education-statistics>

Oates, A., & Johnson, D. (2025). ChatGPT in the classroom: Evaluating its role in fostering critical evaluation skills. *International Journal of Artificial Intelligence in Education*. Advance online publication. <https://doi.org/10.1007/s40593-024-00452-8>

OpenAI. (2025a, April). Introducing OpenAI o3 and o4-mini.

OpenAI. (2025b, April). Thinking with images.

Pai, K.-C., Kuo, B.-C., Liao, C.-H., & Liu, Y.-M. (2021). An application of Chinese dialogue-based intelligent tutoring system in remedial instruction for mathematics learning. *Educational Psychology*, 41(2), 137–152. <https://doi.org/10.1080/01443410.2020.1731427>

- Pardo, A., Jovanovic, J., Dawson, S., Gašević, D., & Mirriahi, N. (2018). Using learning analytics to scale the provision of personalised feedback. *Journal of Learning Analytics*, 5(3), 128–144. <https://doi.org/10.18608/jla.2018.53.15>
- Peschl, M. F. (2024). Human innovation and the creative agency of the world in the age of generative AI. *Possibility Studies & Society*, 2(1), 49–76. <https://doi.org/10.1177/27538699241238049>
- Quantilus. (2024, November). Breaking the AI scaling barrier: How Nvidia and tech leaders are innovating beyond limits. *Quantilus Articles*.
- Renner, M., & Chaban, M. A. V. (2024, April). Real-world gen AI use cases from the world's leading organizations.
- RINF.tech. (2025, July). How multimodal AI is transforming our interaction with technology.
- Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676–688. <https://doi.org/10.1016/j.tics.2016.07.002>
- Sakana AI. (n.d.). Continuous thought machines. Retrieved May 19, 2025, from <https://sakana.ai/ctm/>
- Seneca, L. A. (1917). *Ad Lucilium Epistulae Morales*. Harvard University Press.
- Seton, C., Coutrot, A., Hornberger, M., Spiers, H., Knight, R., & Whyatt, C. (2023). Wayfinding and path integration deficits detected using a virtual reality mobile app in patients with traumatic brain injury. *PLOS ONE*, 18(5), e0282255. <https://doi.org/10.1371/journal.pone.0282255>
- Shahzad, M. F., Xu, S., Lim, W. M., Yang, X., & Khan, Q. R. (2024). Artificial intelligence and social media on academic performance and mental well-being: Student perceptions of

positive impact in the age of smart learning. *Heliyon*, 10(4), e29523.

<https://doi.org/10.1016/j.heliyon.2024.e29523>

Shin, D. (2020). How do users interact with algorithm recommender systems? The interaction of users, algorithms, and performance. *Computers in Human Behavior*, 109, 106344.

<https://doi.org/10.1016/j.chb.2020.106344>

Singh, A. (2024). Empowering data science learners through learnersourcing and student-AI collaboration [Doctoral dissertation, University of Michigan]. University of Michigan Library Deep Blue Repositories. <https://doi.org/10.7302/24958>

Singh, A., Taneja, K., Guan, Z., & Ghosh, A. (2025). Protecting human cognition in the age of AI [Preprint]. arXiv. <http://arxiv.org/abs/2502.12447>

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776–778.

<https://doi.org/10.1126/science.1207745>

Spiers, H. J., Coutrot, A., & Hornberger, M. (2023). Explaining world-wide variation in navigation ability from millions of people: Citizen science project Sea Hero Quest. *Topics in Cognitive Science*, 15(1), 120–138. <https://doi.org/10.1111/tops.12590>

Stanford University Human-Centered Artificial Intelligence Institute. (2025). The 2025 AI Index report.

Sternberg, R. J. (2024). Do not worry that generative AI may compromise human creativity or intelligence in the future: It already has. *Journal of Intelligence*, 12(7), 69.

<https://doi.org/10.3390/jintelligence12070069>

The White House. (2025a, April 23). Advancing artificial intelligence education for American youth (Executive Order).

<https://www.whitehouse.gov/presidential-actions/2025/04/advancing-artificial-intelligence-education-for-american-youth/>

The White House. (2025b, April). Fact sheet: President Donald J. Trump advances AI education for American youth. Retrieved May 19, 2025, from

<https://www.whitehouse.gov/fact-sheets/2025/04/fact-sheet-president-donald-j-trump-advances-ai-education-for-american-youth/>

Time. (2025, April). The definition of multimodal AI.

Topics, E. (2025, March). Number of ChatGPT users (March 2025).

Unite.AI. (2024, July). AI at the International Mathematical Olympiad: How AlphaProof and AlphaGeometry 2 achieved silver-medal standard.

Verbert, K., Duval, E., Klerkx, J., Govaerts, S., & Santos, J. L. (2013). Learning analytics dashboard applications. *American Behavioral Scientist*, 57(10), 1500–1509.

<https://doi.org/10.1177/0002764213479363>

VKTR.com. (2025, March). The scaling of AI foundation models: Progress or plateau?

VKTR.com AI Technology Insights.

Wang, Y., & Zhao, P. (2020). A probe into spoken English recognition in English education based on computer-aided comprehensive analysis. *International Journal of Emerging Technologies in Learning (iJET)*, 15(3), 223–235.

<https://doi.org/10.3991/ijet.v15i03.12937>

<https://doi.org/10.3991/ijet.v15i03.12937>

Ward, A. F. (2013). Supernormal cues: The human brain misinterprets widespread information access as own knowledge. *Psychological Science*, 24(10), 1965–1972.

<https://doi.org/10.1177/0956797613483903>

Ward, A. F. (2021). People mistake the internet's knowledge for their own. *Proceedings of the National Academy of Sciences of the United States of America*, 118(43), e2105061118.

<https://doi.org/10.1073/pnas.2105061118>

Wolbers, T., & Hegarty, M. (2010). What determines spatial reorientation? The role of landmarks and path integration. *Trends in Cognitive Sciences*, 14(3), 138–146.

<https://doi.org/10.1016/j.tics.2010.01.001>

Yavuz, E., Coutrot, A., He, C., McHugh, J., Gürel, M., Hornberger, M., & Spiers, H. J. (2024). Reliance on GPS use is associated with worse spatial memory, but not all types of GPS use may be detrimental. *Acta Psychologica*, 247, 104288.

<https://doi.org/10.1016/j.actpsy.2024.104288>

Zhang, Y., Qin, Y., Chen, Z., Liu, J., Li, Z., Yu, W., & Liu, Y. (2025). A survey of scaling in large language model reasoning [Preprint]. arXiv. <https://arxiv.org/abs/2504.02181>

Appendix A: AI, Academic Pressure, and the Learning Process, a Scenario

Look man, *I get it*. It's Friday, 8 PM. Okay let's be real, it's 11 PM. You've got that problem set due in 59 minutes, and you've been stuck on the first two problems for the whole week. Your dusty Abstract Algebra textbook sits idly by on your desk, your iPad lies askew beside it with a tab open on some forlorn mathematics page in Wikipedia, half the links on the page are blue, the other half are purple—the remnants of the digital rabbit holes you went down, desperately clamoring for some ground to stand on, some concept upon which you can lay the foundation of understanding required to approach this problem. Your laptop spews forth light that fills the room—the sun went to sleep long ago—and your weary eyes drift over to the blinking cursor on a familiar landing page. "What can I help with?" asks the text on screen. It feels like it's mocking you, but at this point, it might be your only way out.

You click the plus button in the bottom right, press on "Add photos and files," and drop in a PDF of your homework.

I need some help with question 2, you type.

Your cursor hovers over the submit icon in a moment of hesitation. You know that this is going to come back to bite you on the next homework and the final. You wonder that maybe if you had just started earlier and put the AI aside, you could have figured this out and avoided making this choice.

But that's the problem: this thing *is* actually helpful. You've written a book together over the past few problem sets, question and answer style. You usually start off saying: "I have a feeling that this particular mathematical object would be useful in solving this problem, but I don't have enough understanding of it to apply it correctly. Can you give me some quick practice problems to prepare me to solve the homework question?" It obligingly spits out a couple

questions, which you dutifully answer and volley back at it, only for it to toss back a graded version that also chidingly shows where you went astray, just as you had prompt engineered it to. You repeat this a couple times feeling like you're actually making some progress towards solving this previously unapproachable problem—but you've got this nagging feeling that you could be doing something else, and a more insidious thought. *Let's be real here, the ChatGPT version of this proof you're working on will get a higher grade than you any day; it's been trained on every Algebra textbook and practice problem there is.*

"I know that," you say aloud, having long ago accepted that as an unfortunate reality. I'm not doing this for the grade, I'm doing it for myself. I'm doing this because I want to learn.

"I still love learning, *right?*"

The silence offers no response.

You decide that you do. *That's why you're here after all.* But you can't help but get this feeling that the whole process has gotten harder lately, that there's somehow more friction between you and solving hard problems.

"What can I help with?" asks the text on screen, again.

You look down at your hands, tapping your thumbs in consideration of the option space.

You type: "Okay, before I even try question 2 again, can you explain the concept of normal subgroups from Chapter 3? I feel like if I don't get that solid, I'm just spinning my wheels. Ask me some clarifying questions as you explain it, so I know I'm actually getting it this time."

The AI interface refreshes, not with a solution, but with a targeted question about the properties of subgroups. A different kind of dialogue begins—slower, more methodical, aimed at foundational understanding rather than immediate answers.

You lean in. It's still late, the problem set is still daunting, but for the first time all week, a flicker of genuine comprehension feels possible. It's not the easy way out, but it feels like a step on the path to actually learning.

You smile, a tired but hopeful one. *What a crazy world we live in.*