Journal of Contemporary Educational Research, 2025, Volume 9, Issue 10
http://ojs.bbwpublisher.com/index.php/JCER
ISSN Online: 2208-8474
ISSN Print: 2208-8466

Study on the Long-Term Cognitive Effects of Excessive LLM Use Among Students and General Users

Bingzhi Li*

School of General Education, Tianjin Foreign Studies University, Tianjin 300204, China

**Author to whom correspondence should be addressed.*

Abstract: Artificial intelligence, especially large language models (LLMs), is reshaping how we learn, communicate, and create. These systems provide students and professionals immediate access to fluent, context-aware language that can support learning, increase productivity, and spark creativity. But what happens when their use becomes excessive? This paper explores the potential long-term consequences of over-relying on LLMs—particularly memory, critical thinking, creativity, and motivation. Drawing on research in cognitive psychology, education, neuroscience, and media studies, it argues that LLMs are best understood as cognitive prostheses: incredibly valuable when used wisely, but risky when they replace rather than support human imagination.

Keywords: Artificial intelligence; Large language models; Cognition; Creativity; Dependency; Education

Online publication: November 10, 2025

1. Introduction

In just a few years, generative artificial intelligence (AI) has gone from a curiosity to a cultural fixture. At the forefront are large language models (LLMs)—systems able to produce text with surprising fluency, wit, and adaptability. Students, teachers, and professionals now turn to them for writing help, translation, tutoring, coding, and brainstorming. In classrooms, especially in ESL contexts, learners from China, Russia, and elsewhere use them for vocabulary support, grammar correction, and even essay drafting. Writers and professionals use them to refine tone, outline ideas, or speed up their workflow.

Educators often share the same initial reaction: LLMs seem like shortcuts, even cheats, undermining genuine learning. Yet many soon recognize their unexpected sophistication. These models do not just regurgitate—they generate context-sensitive phrasing, nuanced vocabulary, and coherent arguments. For language learners, access to polished expression that once required years of immersion can feel nothing short of miraculous.

But every miracle has a shadow. At what point does helpful scaffolding slip into harmful substitution?

When does reliance turn into dependency? And most crucially: does heavy use of generative AI blunt people's own ability to think and create? These questions strike at the heart of cognitive development across the lifespan [1,2]. Human thinking has always been shaped by its tools—and if LLMs increasingly mediate our thought, the long-term consequences deserve careful scrutiny.

This paper examines those consequences. It reviews existing research on digital technology and cognition, builds a theoretical framework combining scaffolding, cognitive load theory, the extended mind, and media ecology, and weighs the benefits and risks of LLM use across different contexts. It closes with practical recommendations to ensure these tools support, rather than supplant, human cognitive growth.

2. Literature review

2.1. Digital technology and cognition

Every major communication tool has reshaped how people think. The printing press spread literacy but changed memory practices. The calculator streamlined arithmetic yet provoked debates about fading numeracy. The internet gave us instant access to knowledge, but encouraged what researchers call "cognitive offloading" or the "Google effect"—we remember less when we know information can be looked up [3].

LLMs push this trajectory further. Unlike search engines that retrieve, they generate. They do not just hand back facts; they compose arguments, explanations, and even creative passages. This raises a sharper question: when we outsource ideation, drafting, and sense-making itself, how does that reshape the internal processes of memory, reasoning, and creativity?

2.2. Cognitive development and learning theories

Classical theories of learning stress the importance of struggle and active engagement in building real understanding. Piaget [4] argued that growth happens when learners face challenges that unsettle their current ways of thinking, forcing them to adapt and reorganize. Vygotsky [5] took this further by introducing the idea of the zone of proximal development (ZPD)—the sweet spot where learners, with the right kind of support, can stretch just beyond what they can currently do. But crucially, support is meant to fade. Scaffolding only works if it is temporary, helping learners internalize strategies so they eventually stand on their own.

Cognitive load theory [6] adds another dimension: our working memory is limited. Good instruction should lighten unnecessary burdens without removing the mental effort needed for deeper learning. LLMs, for example, can ease tasks like grammar checking or formatting, clearing space for higher-order thinking. But if they also take over the heavy lifting—such as structuring arguments or generating insights—they may prevent learners from forming the mental schemas that underpin long-term understanding.

2.3. Creativity and originality

Creativity thrives on openness, autonomy, and the ability to produce ideas that are both novel and useful [7]. It often emerges from messy cycles of trial and error, critique, failure, and revision [8]. Generative AI can mimic parts of this cycle at incredible speed, offering a flood of prompts and drafts. This can be energizing, but it also risks short-circuiting the slower processes—incubation, persistence, and frustration—that often lead to breakthroughs. Moreover, what these systems produce is not truly original: their outputs are a recombination of existing data, shaped by their training sets and design biases [9]. Their creativity, in other words, has boundaries that human imagination does not.

2.4. Compulsive technology use

Many digital platforms are built to claim and hold our attention, rewarding users with constant stimuli that form habits [10]. LLMs, with their conversational ease and instant assistance, can foster similar patterns. For students under pressure or people seeking reassurance, it is easy to slip into checking them repeatedly. The danger here is subtle: if answers are always available, tolerance for ambiguity and patience for struggle shrink. Intellectual discomfort—the fuel of growth—may be replaced by the comfort of constant support.

2.5. Neuroscience of learning and media interaction

The brain adapts to practice. Skills we use often strengthen related neural circuits [11]. Skills we outsource, by contrast, can weaken. For example, reliance on GPS correlates with less activation in hippocampal regions that handle spatial memory [12]. By analogy, if learners routinely hand off ideation and composition to LLMs, neural circuits that support memory retrieval, working memory rehearsal, and associative creativity could be underused over time. This remains a hypothesis: research on LLMs' neurocognitive effects is still in its infancy, and long-term studies are urgently needed.

2.6. Motivation and self-determination

Deci and Ryan's self-determination theory [13] suggests that intrinsic motivation comes from feeling autonomous, competent, and connected. Mastery fuels persistence and curiosity. When LLMs provide ready-made answers, they can deliver quick wins—better grades, faster results—but also rob learners of the satisfaction of overcoming difficulty. This shortcut effect may be especially problematic in early educational stages, when habits of discipline, problem-solving, and self-reflection are still being formed.

3. Theoretical framework

This study synthesizes several theoretical lenses to analyze ChatGPT's cognitive implications: Vygotskian scaffolding, cognitive load theory, the extended mind thesis [14], media ecology [15], and distributed cognition [16]. Together, these frameworks allow us to examine ChatGPT both as a pedagogical scaffold and as an environmental affordance that alters the distribution of cognitive labor between agents and artefacts. They also encourage attention to the sociocultural contexts in which AI is deployed, recognizing that effects are mediated by instructional design, institutional incentives, and cultural norms.

3.1. Vygotsky's scaffolding and the zone of proximal development

Through a Vygotskian lens, large language models can act like a "more knowledgeable other," stretching a learner's capacity by offering support within their ZPD. When teachers use these tools thoughtfully—posing gradually harder prompts and encouraging reflection—they can foster deep learning. But if students or instructors lean on them as a permanent crutch, letting the machine do the reasoning, the scaffolding collapses into dependency. The key is design: pairing AI help with reflective questions, transfer exercises, and tasks that grow progressively more independent helps preserve the integrity of the ZPD.

3.2. Cognitive load and generative AI

Cognitive load theory helps explain both the promise and the risk of LLMs. Novices often benefit from reduced "extraneous load"—having the system handle tasks like grammar or formatting frees up space for real learning. But for more advanced learners, overuse becomes harmful if the model also does the conceptual heavy lifting,

such as crafting arguments or synthesizing ideas. Effective integration requires tailoring support to the learner's stage: more procedural assistance for beginners, more insistence on independent synthesis for those further along.

3.3. The extended mind and distributed cognition

The "extended mind" thesis suggests our thinking does not stop at the skull—it can stretch across tools and environments. LLMs might, in theory, become part of this extended system. Yet there is a catch: unlike a calculator or a peer collaborator, these models are black boxes. Their training data, probabilistic reasoning, and occasional errors are hidden from the user. Distributed cognition works best when partners are transparent and accountable; when one partner is not, the result may be brittle thinking—appearing fluent on the surface but lacking depth or scrutiny.

3.4. Media ecology and habits of thought

From a media ecology perspective, LLMs are not just tools but environments that shape how we think. Their conversational rhythm encourages speed and polish, often skipping over the messy drafting stages where ideas evolve. Over time, classrooms and cultures may adapt to these norms, producing writers and thinkers who expect instant refinement. This shift is not necessarily bad—but it is significant. Traditions of originality, slow reflection, and rhetorical struggle may give way to smoother, but potentially shallower, forms of expression.

4. Methodological considerations and empirical approaches

To move past speculation, researchers should design empirical studies that capture the real cognitive effects of sustained AI use. A few promising approaches include:

(1) Longitudinal cohort studies: Follow students at different developmental stages (late childhood, adolescence, early adulthood) over multiple years. Track their AI use, academic outcomes, memory, creativity, and motivation, combining standardized tests with interviews about their sense of voice and authorship.

(2) Experimental interventions: Run randomized trials where groups receive different levels of AI support—none, limited procedural help, or full generative assistance. Measure immediate performance, retention after delays, and the ability to transfer skills to new situations.

(3) Neurocognitive and neuroimaging studies: Use tools like fMRI and EEG to compare habitual AI users with non-users. Focus on regions tied to executive function, memory consolidation, and creative association, using both cross-sectional and long-term designs to capture changes over time.

(4) Mixed-methods ethnography: Pair data with classroom and workplace observations. Ethnographic research can reveal how local cultures, assessment pressures, and institutional norms shape whether AI is used as a scaffold for thinking or simply as a shortcut.

(5) Case studies and vignettes: Develop detailed portraits of learners who either thrive with AI support or fall into dependency. Such stories highlight the role of teacher guidance, assignment design (see **Appendix**), and individual differences in self-regulation.

4.1. Illustrative vignettes

To see how this plays out in practice, consider the following contrasting examples:

Vignette 1: The ESL learner who leverages AI wisely

Mei, a university student learning English, uses an LLM to test idiomatic phrases, check complex

grammar, and explore richer vocabulary. But she does not stop there. She annotates her drafts, explaining why she accepts some AI suggestions and rejects others. Over a semester, her fluency grows, her awareness of language deepens, and her confidence improves. For Mei, the AI functions as a scaffold—support that fades as her own competence strengthens.

Vignette 2: The student who becomes dependent

Alex, by contrast, uses AI to draft nearly all of his written assignments. He edits very little, relying on polished outputs for quick submission. While his grades rise in the short term, his performance falters in timed essays where AI cannot assist. His writing shows fluency without depth. Alex's case illustrates how easy substitution can produce immediate results but undermine long-term skill development.

Vignette 3: The professional collaborator

Fatima, an urban planner, uses an LLM to draft early versions of policy documents and to generate illustrative scenarios. These drafts save her hours, freeing time to engage more directly with experts and community members. Yet she draws a clear line: for final public-facing documents, she scales back AI use, ensuring that human oversight and community voices remain central. Fatima's approach shows how AI can be integrated appropriately—supporting efficiency without displacing professional judgment or accountability.

4.2. Practical classroom interventions

Teachers do not need to choose between banning AI and embracing it uncritically. Instead, they can adopt strategies that preserve cognitive growth while taking advantage of the technology's benefits:

(1) Shifting from product to process: Ask students to submit drafts in stages, annotate their AI interactions, and reflect on why they accepted or rejected certain suggestions. This rewards metacognitive thinking, not just polished final products.

(2) Teaching source evaluation through AI: Use AI-generated text as a teaching tool. By asking students to identify subtle errors or biases in machine outputs, instructors can train critical thinking in a low-stakes environment.

(3) Building AI-free practice time: Reserve portions of class for handwritten or typed work without AI assistance. Alternating between "with AI" and "without AI" modes helps students maintain flexibility and resilience in their own thinking.

(4) Promoting collaborative AI use: Encourage group projects where students must debate, critique, and adapt AI outputs together. This not only deepens processing but also creates shared responsibility for quality.

5. Creativity as collaboration: Benefits and risks of LLM integration

5.1. Benefits of generative AI

Despite the risks, the advantages of thoughtful use are clear.

Language learning and communication: For learners of a second language, LLMs can be transformative. They offer tailored corrections, natural phrasing, and exposure to different registers and dialects. When students are asked to explain, translate, or build on the AI's feedback, the tool becomes a scaffold that accelerates fluency instead of replacing learning.

Accessibility, inclusion, and equity: LLMs also lower barriers. They offer affordable tutoring, instant writing support, and communication aids that can help students with fewer resources or those with disabilities. This democratizing potential is worth protecting—though always with an eye toward avoiding new forms of depen-

dency.

Creativity as collaboration: When treated as a partner rather than a shortcut, large language models can broaden creative possibilities. They might suggest fresh metaphors, sketch out structural outlines, or offer vocabulary that sparks new directions. The real value emerges in collaboration: when human judgment, context, and values guide and refine the model's suggestions, the outcome can be something neither could achieve alone.

Efficiency and expert augmentation: In professional settings, LLMs shine at handling repetitive tasks—drafting boilerplate, scaffolding code, or synthesizing long documents. By offloading these routine burdens, experts can spend more energy on interpretation, strategy, and ethical decision-making. In this way, the technology serves less as a replacement than as an amplifier, boosting productivity without necessarily diminishing expertise.

5.2. Risks of excessive or compulsive use

Still, alongside these benefits come real risks, especially when use becomes habitual or unchecked. Educators, professionals, and policymakers need to take these challenges seriously.

Cognitive offloading and memory: Research on the "Google effect" shows that when people know information can be stored externally, they encode and retain less of it themselves [3]. The same pattern may occur with LLMs: when explanations, arguments, or full drafts are routinely outsourced, learners lose opportunities for rehearsal and elaboration—the very processes that strengthen memory. Over time, this could mean weaker recall of facts, arguments, and rhetorical strategies that once would have been internalized.

Critical thinking and vigilance: AI-generated text is often fluent and persuasive, but not always accurate. Outputs may contain subtle errors, omissions, or cultural biases [9]. Relying on these outputs without critical engagement can erode skills in evaluating credibility, testing logic, and spotting fallacies. The result is a fragile epistemic culture, where people accept polished answers without probing their foundations.

Creative habits at risk: True creativity often demands persistence through uncertainty, awkward drafts, and failed attempts [8]. When polished alternatives are available instantly, it becomes tempting to bypass that struggle. Over time, this can reduce tolerance for ambiguity, weaken resilience, and narrow the range of individual and cultural expression.

Shifts in motivation and identity: Outsourcing intellectual effort can also reshape how people see themselves. A student who habitually relies on AI for essays may begin to credit success to the tool, not their own abilities. Professionals who lean on AI for phrasing may feel their personal voice slipping away. Conversely, the crutch provided by AI may mislead the user into assuming credit for authorship unearned and, subliminally, into over-estimating their own ability. These shifts matter: they can undermine intrinsic motivation, dilute creative identity, and raise thorny questions about authorship and integrity. Plagiarism, always an issue, is thus given an expanded definition.

6. Discussion

6.1. Compulsive use and mental health concerns

Because LLMs respond instantly and feel conversational, they can encourage patterns of constant checking and reliance, similar to other digital compulsions [10]. For vulnerable groups—such as adolescents or people with anxiety disorders—this reliance can deepen existing struggles. It may fuel procrastination, avoidance of difficult tasks, and heightened stress tied to performance expectations. These psychological consequences deserve close

attention and monitoring.

6.2. Cross-cultural and developmental considerations

The impact of LLMs is not the same everywhere. How they are used depends heavily on cultural views of authority, authorship, and education. In systems that prize rote memorization and high-stakes testing, these tools may be used mainly for efficiency—polishing submissions rather than supporting genuine learning—which raises the risk of substitution. In contrast, educational cultures that value inquiry and project-based learning are more likely to use them in ways that preserve independent thinking.

Age also matters. Younger learners are especially vulnerable because their cognitive, motivational, and self-regulatory skills are still developing. Adolescents, whose executive functions are still maturing [17], may be particularly prone to impulsive over-reliance. Adults, by comparison, often have stronger meta-cognitive habits and may integrate these tools more strategically. This suggests the need for differentiated approaches: tighter safeguards and guided use for younger students, and reflective, incentive-based practices for older learners.

6.3. Ethical and institutional concerns

Generative AI raises pressing questions about authorship, plagiarism, and academic honesty. Institutions must adapt by revising honor codes, redesigning assessments, and rethinking evaluation practices. Yet framing all AI use as "cheating" risks losing valuable opportunities to teach responsible engagement. A more balanced approach would require transparency—students openly declaring their use of AI—paired with assessments that reveal authentic learning. Examples include oral defenses, process portfolios, or in-class writing tasks that ensure internalized competence.

6.4. Policy and pedagogical recommendations

To balance benefits and risks, educators and institutions should pursue multi-layered strategies:

(1) Curriculum design for AI-aware learning: Build AI literacy into courses. Teach students about biases, limitations, and the distinction between generating ideas and verifying them. Assignments should include process documentation, such as draft logs and reflections, to highlight iterative development.

(2) Scaffolded AI use and withdrawal: Begin by allowing AI to reduce surface-level demands, then gradually withdraw support as students demonstrate competence. Assessments should increasingly require independent application of skills.

(3) Assessment reform: Move beyond take-home essays as the sole metric. Incorporate oral exams, real-time problem-solving, and portfolios. Where AI is used, students are required to declare and reflect on its role in their process.

(4) Teaching critical evaluation skills: Explicitly train students to fact-check, cross-reference, and spot AI-generated errors. Case studies of flawed outputs can help build skepticism and verification habits.

(5) Protecting younger learners: Limit unsupervised use of AI for high-stakes work. Emphasize handwritten drafts, peer workshops, and other activities that preserve the value of struggle and reflection.

(6) Institutional policies and support: Universities and schools should issue clear guidelines for responsible use, train faculty on assignment design, and give students resources for AI literacy. Policymakers should also support long-term research into cognitive effects and ensure equitable access to beneficial tools.

7. Limitations and directions for further research

At present, our understanding of the long-term cognitive impact of LLMs is still very limited. Most existing studies look only at short-term effects—how people perform on immediate tasks or their initial experiences with the tools. What we lack are longitudinal studies that follow learners over many years to see whether consistent reliance on AI weakens memory, critical thinking, or creative fluency—and whether any losses can be reversed.

Neuroscience also has a role to play. Imaging studies could explore whether habitual outsourcing of writing or problem-solving to AI changes patterns of brain activity in areas linked to memory retrieval, working memory, or creative association. Alongside this, carefully designed experiments should test classroom strategies that encourage reflective, balanced use of AI, and measure how effective they are across different ages and cultural contexts.

This paper has been primarily theoretical, drawing connections across digital media studies, cognitive psychology, and educational theory to suggest possible outcomes of sustained LLM use. Because the technology is so new, empirical research specific to these models is still limited. Causal claims about neurocognitive change, for instance, require long-term behavioral and neuroimaging studies that do not yet exist. In addition, cultural differences in pedagogy, labor markets, and infrastructure mean that findings from one setting may not apply universally. Future research must therefore be cross-cultural, multilingual, and sensitive to local conditions.

7.1. Expanding the evidence base: Key questions for future research

To move from speculation to evidence, researchers need to focus on several pressing questions:

Dose-response relationships: How much and how often can people use AI before measurable cognitive changes appear? Are there thresholds beyond which the effects become lasting?

Reversibility: If overreliance dulls certain skills, can deliberate training or scaffolded withdrawal bring them back?

Moderators: How do factors like working memory capacity, metacognitive skills, or cultural attitudes shape outcomes?

Domain specificity: Are some areas of knowledge—such as math problem solving, narrative writing, or scientific reasoning—more vulnerable to outsourcing than others?

Equity implications: Do AI tools help close learning gaps by widening access, or do they deepen divides by benefiting strategic users while leaving others dependent?

7.2. Policy considerations for governments and accrediting bodies

Policy responses must balance innovation with safeguards for development. National ministries and accrediting organizations could:

Update curricula to include AI literacy standards and new assessment frameworks.

Fund professional development for teachers to help them integrate AI thoughtfully.

Promote transparency in AI use and support open-access platforms for under-resourced schools.

Establish accreditation standards that ensure graduates master core competencies, not just the ability to generate polished outputs with AI.

7.3. Industry responsibility and product design

AI developers shape not only the tools but also the learning environments in which they are used. Thoughtful design can reduce risks by:

Offering an “assist mode” that gradually reduces support, encouraging independence.

Providing provenance and confidence scores to help users judge reliability.

Ensuring affordable, privacy-protective options for schools and universities.

Partnering with educators to design tools and curricula that build AI literacy and critical thinking.

7.4. Societal and labor market considerations

The ripple effects of overreliance on generative AI extend beyond the classroom. In fields that value originality—like journalism, creative writing, or research—heavy dependence could push output toward uniformity, privileging speed and clarity over nuance and novelty. Yet in other areas, where routine synthesis consumes time and offers little intellectual payoff, AI can free professionals to focus on interpretation, strategy, and relational work. Policymakers and industry leaders should therefore avoid one-size-fits-all rules, instead creating domain-specific guidelines that balance efficiency gains with the protection of expertise.

7.5. Ethical reflections: Agency, authorship, and cultural memory

Beyond practical concerns lies a deeper set of questions about what kind of culture we want to build. AI tools that generate language at scale inevitably shape what gets preserved and valued in public discourse. If machine-mediated phrasing becomes the norm, whose voices will be amplified, and whose will be pushed aside? Authorship is not just a technical issue—it touches accountability, recognition, and the distribution of cultural capital. Academic institutions and cultural producers must therefore grapple with not only the "how" of AI integration but also the "why": what do we want to honor and protect in human creativity and intellectual labor?

7.6. A note on pedagogical optimism

Although much of this paper has focused on risks, there are strong reasons for optimism. Educators have always adapted to disruptive technologies by rethinking how they teach [18]. The same is true here: the capabilities of large language models can be used to enrich formative learning, personalize feedback at scale, and design assessments that target higher-order skills rather than rote production. Moving forward will require careful design, ongoing research, and collaboration across education, industry, and policy—but the potential for positive transformation is real.

8. Conclusion and implications

Generative AI is undeniably transformative and offers enormous potential. Used with care, it can scaffold learning, widen access, amplify productivity, and inspire creativity in ways that once seemed unimaginable. Students can experiment with new styles of writing, professionals can streamline complex workflows, and lifelong learners can receive personalized support on demand. Yet, as the history of technology reminds us, the cognitive effects of innovation are never neutral. The very qualities that make it appealing—its fluency, immediacy, and scaffolding power—also carry risks when overused. Too much reliance can blunt memory, weaken critical judgment, narrow creative habits, and undermine intrinsic motivation.

The long-term effects will depend not on the technology alone but on how we design its use: through pedagogy, institutional policies, and cultural norms. The answer is not to reject AI outright but to build environments—curricula, assessments, and practices—that protect human agency, nurture reflection, and ensure AI acts as a tool for augmentation rather than replacement.

If human effort is displaced entirely—if drafting, ideation, and problem-solving are handed over wholesale to machines—the risks become clear. Memory may weaken because information is never rehearsed or retrieved.

Critical thinking may be dulled because arguments are accepted at face value rather than interrogated. Creativity may narrow because the messy process of trial and error is replaced with instant, polished alternatives. What we stand to lose are not just isolated skills but the resilience, patience, and reflective habits that underpin lifelong learning and authentic innovation.

The antidote lies in design. Curricula must deliberately weave AI into learning while protecting opportunities for independent struggle and original thought. Assessments should look beyond the finished product to examine the process—drafts, reflections, and revisions that reveal how students are engaging with ideas. Policies at institutional and national levels must strike a balance: permitting AI where it scaffolds growth, restricting it where it undermines mastery, and always insisting on transparency in its use.

When approached in this way, generative AI can amplify rather than diminish human capacities. It can reduce unnecessary friction while preserving the essential challenges that make learning durable and creativity meaningful. The goal is not to strip away the struggle but to ensure that it remains purposeful—allowing AI to handle the routine so that humans can invest energy in higher-order thinking, deep reflection, and imaginative exploration. Done wisely, this integration can help create an intellectual culture in which technology extends rather than erodes what is most valuable in human cognition.

8.1. Bringing the threads together: Practical epilogue

The central claim of this paper is that generative AI is neither wholly good nor wholly harmful. It is better understood as a powerful cognitive tool whose effects are determined by how, when, and why it is used. In this sense, the metaphor of a cognitive prosthesis is especially instructive. Just as a prosthetic limb can restore or even enhance physical capacity, AI can extend intellectual capacity—helping learners draft, organize, or generate ideas that might otherwise be out of reach. But prostheses come with trade-offs. If someone relies on a prosthetic device without maintaining their underlying muscles, those muscles may atrophy. In the same way, if learners or professionals lean exclusively on AI to handle thinking tasks, the very skills AI is meant to support—memory, critical reasoning, creative problem-solving—may weaken over time.

This tension highlights the need for what might be called rehabilitation-style pedagogy. Education systems should not only permit AI use but deliberately design practice routines that strengthen internal cognitive abilities alongside it. For example, students might use AI for early brainstorming but be required to produce final drafts independently. Or learners could analyze and critique AI-generated outputs to sharpen their evaluative judgment, rather than simply submitting them as is. In this model, AI is not a replacement for effort but a structured partner in developing resilience, reflection, and skill.

The challenge, then, is balance. Just as physical therapy integrates prosthetics into daily life while keeping the body active, classrooms and workplaces need to integrate AI in ways that expand possibilities without hollowing out capability. Done thoughtfully, this approach can preserve the formative struggles—moments of confusion, revision, and perseverance—that make learning durable and creativity authentic, even as AI lightens some of the burdens along the way.

8.2. Reflections on human flourishing in an AI-augmented age

Beyond memory, reasoning, and creativity lies a larger question: what does it mean to live a flourishing intellectual life when many cognitive tasks can be outsourced? History offers precedents—mechanized farming freed people from manual labor, calculators reshaped numeracy, and the internet changed how we remember facts. The challenge now is to ensure that AI enhances, rather than erodes, the human experience. Done thoughtful-

ly, AI can foster reflection, free time for higher pursuits, and widen access to knowledge—without erasing the struggles and satisfactions that give learning and achievement their meaning.

8.3. Closing call to scholars, educators, and policymakers

This paper closes with a call to action. Researchers need to move beyond short-term observations and build rigorous longitudinal and experimental evidence on how LLMs affect cognition. Educators must design classrooms that preserve genuine learning, spaces where students practice metacognition, reflection, and independent thinking, even as they benefit from AI support. Institutions should craft balanced policies that encourage transparency and equitable access, rather than defaulting to bans that ignore reality. And industry leaders must take responsibility by creating tools that support, rather than undermine, pedagogy—features that guide users toward critical engagement and reflective use.

The future of human cognition will not be written by algorithms alone. It will be shaped by our collective choices about pedagogy, policy, and professional norms. If we act thoughtfully and collaboratively, we can harness the benefits of AI while avoiding the trap of substitution—protecting memory, critical thought, and creativity as essential cornerstones of intellectual life.

Disclosure statement

The author declares no conflict of interest.

References

[1] Luria AR, 1973, The Working Brain: An Introduction to Neuropsychology, Basic Books.

[2] Sternberg RJ, 2019, Cognitive Psychology (7th ed.), Cengage Learning.

[3] Sparrow B, Liu J, Wegner DM, 2011, Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips. Science, 333(6043): 776–778.

[4] Piaget J, Inhelder B, 1966, The Psychology of the Child, Basic Books, 72–75.

[5] Vygotsky LS, 1978, Mind in Society: The Development of Higher Psychological Processes, Harvard University Press.

[6] Sweller J, Ayres P, Kalyuga S, 2011, Cognitive Load Theory, Springer, New York.

[7] Runco MA, Jaeger GJ, 2012, The Standard Definition of Creativity. Creativity Research Journal, 24(1): 92–96.

[8] Sawyer RK, 2019, Explaining Creativity: The Science of Human Innovation (2nd ed.), Oxford University Press.

[9] Bender EM, Gebru T, McMillan-Major A, et al., 2021, On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, 610–623.

[10] Alter A, 2017, Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked, Penguin Press.

[11] Draganski B, Gaser C, Busch V, et al., 2004, Neuroplasticity: Changes in Grey Matter Induced by Training. Nature, 427(6972): 311–312.

[12] Maguire EA, Gadian DG, Johnsrude IS, et al., 2000, Navigation-Related Structural Change in the Hippocampi of Taxi Drivers. Proceedings of the National Academy of Sciences, 97(8): 4398–4403.

[13] Deci EL, Ryan RM, 2000, The "What" and "Why" of Goal Pursuits: Human Needs and the Self-Determination of Behavior. Psychological Inquiry, 11(4): 227–268.

[14] Clark A, Chalmers D, 1998, The extended mind. Analysis, 58(1): 7–19.
[15] McLuhan M, 1964, Understanding Media: The Extensions of Man, McGraw-Hill.
[16] Hutchins E, 1995, Cognition in the Wild, MIT Press.
[17] Casey BJ, Jones RM, Somerville LH, 2011, Braking and Accelerating of the Adolescent Brain. Journal of Research on Adolescence, 21(1): 21–33.
[18] Mishra P, Koehler MJ, 2006, Technological Pedagogical Content Knowledge: A Framework for Teacher Knowledge. Teachers College Record, 108(6): 1017–1054.

Appendix

Sample Assignment Designs

Assignment 1: Scaffolded Research Essay

Week 1: Students explore a topic without AI, submitting notes and research questions.

Week 2: Students may consult an LLM for framing questions and outline suggestions, but must annotate each AI interaction and explain their choices.

Week 3: Drafts include a process log, annotated AI transcripts, and a reflection on how their reasoning evolved.

Week 4: Students defend their argument orally to demonstrate independent understanding.

Assignment 2: Creative Writing Workshop with AI Prompts

Pre-work: Students develop their own story premise without AI.

Workshop: Peers exchange feedback. Students then use AI prompts (e.g., shift point of view, alter the ending, change tense) in three rounds, documenting their choices.

Final: Students submit the story along with a critical reflection on which ideas came from them, which came from AI, and how they synthesized the two.

These examples show how AI can be used as a scaffold while ensuring that core processes of learning and creativity remain in the hands of the student.