Clareus Scientific Science and Engineering
Volume 2 Issue 10 December 2025
DOI: 10.70012/CSSE.02.062
ISSN: 3065-1182

Coding with ChatGPT: Empirical Evidence of Cognitive Offloading in Computer Science Education

Citation: Robin Vivian. "Coding with ChatGPT: Empirical Evidence of Cognitive Offloading in Computer Science Education". Clareus Scientific Science and Engineering 2.10 (2025): 10-19.

Article Type: Research Article

Received: October 15, 2025

Published: December 02, 2025

Robin Vivian*

Laboratoire Perseus, University of Lorraine, Metz, France

***Corresponding Author:** Robin Vivian, Laboratoire Perseus, University of Lorraine, Metz, France.

Abstract

Generative Artificial Intelligence (AI) tools such as ChatGPT, Mistral, and Copilot are reshaping the educational landscape, particularly in programming and computer science education. Their capacity to generate, debug, and optimize code provides immediate performance benefits, yet their influence on long-term cognitive development remains uncertain. This article merges empirical evidence from a quasi-experimental study conducted with 151 first-year computer science students and critical reflection on the pedagogical foundations that have guided education for over half a century. Results indicate significant short-term performance gains in AI-assisted tasks (20-40% improvement) but weak correlations with unaided problem-solving performance ($r \approx 0.15$, $p > 0.05$), suggesting limited long-term learning transfer. Drawing on Bloom's taxonomy and cognitive offloading theory, this paper explores how reliance on AI may alter students' metacognitive processes, diminish their ability to think algorithmically, and challenge the traditional models of knowledge acquisition. Beyond empirical findings, it argues for a new pedagogy of AI — one that balances digital assistance with critical reflection, autonomy, and ethical literacy. Generative AI should be viewed not as a cognitive shortcut but as a partner in co-creation, requiring structured integration to prevent the emergence of "artificial learners" with diminished cognitive depth.

Introduction

"*I have heard, said Socrates, that near Naucratis in Egypt there was one of the ancient gods of that country to whom the Egyptians dedicated the bird they call the ibis; this demon is called Theuth; it was he who invented numeration and calculation, geometry and astronomy, backgammon and dice, and finally writing. Theuth came to King Thamus and showed him the arts he had invented, saying that they should be spread among the Egyptians. 'The teaching of writing, O King,' said Theuth, 'will increase the knowledge and memory of the Egyptians, for I have found the remedy for forgetfulness and ignorance*".

The king replied: "Ingenious Theuth, one is capable of creating the arts, another of judging to what extent they will harm or benefit those who must use them: thus you, father of writing, voluntarily attribute to it an effectiveness contrary to that of which it is capable; you have found a way, not to retain but to renew memory, and what you will give your disciples is the presumption that they have knowledge, not *knowledge itself; for when they have read much without learning, they will believe themselves to be very learned, and they will most often be nothing but ignorant people of troublesome company, because they will believe themselves to be learned without being so*".

We are making the same errors in judgement in 2025.

The Promise and Peril of Generative Artificial Intelligence

Generative AI tools have rapidly transitioned from novelty to necessity in higher education. With the emergence of large language models (LLMs) such as GPT-4, Mistral, and Gemini, students now access systems capable of composing essays, designing algorithms, or producing entire applications in seconds. In programming education, where debugging and logical reasoning were once central to learning, AI can now generate correct code from a simple prompt. These technologies undeniably improve productivity and engagement, offering real-time feedback and personalized explanations. Yet, this convenience raises an essential pedagogical question: ***Does AI enhance learning, or merely simulate it?***

AI's pedagogical potential is well documented. Studies have demonstrated improvements in immediate task performance, motivation, and accessibility when AI is integrated into coursework (Zawacki-Richter et al., 2023; Peñalvo, F. J. 2023). Adaptive learning systems can tailor feedback to individual students, helping them overcome difficulties in real time. However, concerns have emerged that such tools may also foster ***cognitive dependency*** — a reliance on external intelligence that replaces deep reasoning with instant retrieval (Shalu et al. 2025).The paradox of generative AI in education lies precisely here: it amplifies performance while potentially eroding the very skills that education seeks to cultivate — reflection, creativity, and independent thought.

Historical Context: From Bloom to the Machine

Since Bloom's *Taxonomy of Educational Objectives* (1956), educators have sought to scaffold learning through progressive cognitive stages — from remembering and understanding to analyzing, evaluating, and creating. This model, revised by Anderson and Krathwohl (2001), remains central in teacher training and curriculum design. Yet, the arrival of generative AI calls into question the viability of this framework. When a student can generate a fully functional algorithm or essay with a single query, what becomes of the intermediate cognitive steps — comprehension, analysis, and synthesis — once considered essential to deep learning?

Recent research (Slimi, 2023; Gonsalves C., 2024) suggests that LLMs disrupt Bloom's hierarchy by enabling learners to bypass these stages entirely. The act of *creating* becomes a simulation — the student no longer constructs knowledge but curates AI output. This phenomenon aligns with theories of ***cognitive offloading*** (Risko & Gilbert, 2016), where individuals delegate mental processes to external systems. While offloading can free cognitive resources for higher-order reasoning, excessive reliance on AI risks atrophying the very abilities that define intellectual autonomy.

Aims and Scope

This paper aims to reconcile two perspectives often treated separately:

1. The ***empirical dimension***, examining measurable effects of AI on learning outcomes in programming education; and
2. The ***conceptual dimension***, exploring how AI reshapes foundational cognitive models such as Bloom's taxonomy and metacognitive learning.

Building on an experimental study of 151 undergraduate computer science students, the article assesses short-term and long-term impacts of AI integration in assessment contexts. It combines quantitative analyses (exam performance, ANOVA, correlation measures) with theoretical discussion on dependency, ethics, and pedagogy. Beyond describing performance differentials, this work interrogates the philosophical and educational implications of a world where ***machines not only assist learning but perform it***.

Research Questions

This integrated study seeks to answer four central questions:

1. To what extent does generative AI improve short-term student performance in programming tasks?
2. Does this improvement translate into durable learning and skill retention?
3. How does the pervasive use of AI challenge traditional pedagogical frameworks such as Bloom's taxonomy?
4. What forms of AI literacy and hybrid pedagogy could sustain human cognitive development in the age of automation?

Theoretical Framework

Learning Hierarchies in the Age of Artificial Intelligence

For decades, Bloom's taxonomy has structured pedagogical objectives through a hierarchical model: remembering, understanding, applying, analyzing, evaluating, and creating (Bloom et al., 1956; Anderson & Krathwohl, 2001). Its premise is that meaningful learning results from progressively internalizing knowledge through increasingly complex cognitive operations. Yet, the rise of generative AI tools destabilizes this sequence. When a student can instantly obtain an accurate explanation or working solution, the mental trajectory from comprehension to creation collapses into a single act of *prompting*.

Generative systems thus challenge not only *how* we learn, but *what* it means to know. They perform the acts once used to evidence mastery — analysis, synthesis, evaluation — transforming them into automated processes. As Slimi (2023) argues, the risk is no longer ignorance but ***illusion of competence***: a learner may believe they understand because the system provides coherent output. This illusion mirrors the "Google effect" on memory (Sparrow et al., 2011), now extended to reasoning itself.

However, AI's role needs not be destructive. Cognitive offloading — delegating low-level cognitive tasks to external systems — can be beneficial if it frees mental resources for higher-order reasoning (Risko & Gilbert, 2016). The key distinction lies in ***intentionality***: offloading as a strategy for reflection versus offloading as avoidance of thought. Education's challenge is to cultivate *metacognitive regulation* — helping students decide when and why to rely on AI, and how to internalize what it produces.

Reinterpreting Bloom's Domains under Generative AI

Bloom's model extends beyond the cognitive domain to include affective and psychomotor aspects, emphasizing attitudes, values, and motor skills (Krathwohl et al., 1964). Generative AI exerts pressure on all three.

1. ***Cognitive Domain***: AI systems transform the processes of memorization, comprehension, and application. Knowledge retrieval becomes instantaneous; comprehension risks becoming superficial; application is often replaced by imitation. As demonstrated in LLM-supported programming education, students can now bypass reasoning stages by querying an AI for debugging or algorithm design — achieving correct outputs without internalized understanding. This trend is already evident in the industrial sector, particularly in the recruitment of programmers (Chen 2025).
2. ***Affective Domain***: Engagement, curiosity, and persistence — once central to learning motivation — are undermined when AI provides immediate answers. Students may experience gratification without effort, reducing the emotional investment traditionally associated with mastery (Deci & Ryan, 2020).
3. ***Psychomotor Domain***: While less relevant to theoretical subjects, the domain still applies to digital skills. The integration of AI into integrated development environments (e.g., Copilot, Replit) reduces manual coding effort, transforming programming from an *embodied craft* into a curatorial act of prompting and reviewing.

Thus, generative AI introduces a ***post-taxonomic paradigm***, where learning becomes distributed between human and machine cognition. As Annuš N. (2024) suggests, this calls for redefining educational taxonomies not around hierarchical levels but around *collaborative intelligences* — human, artificial, and hybrid.

The Paradox of Educational Automation

AI's integration into education echoes earlier technological waves — from calculators to search engines — yet its generative capacity introduces a qualitative rupture. Whereas prior tools accelerated existing cognitive processes, LLMs *replace* them. A student who once struggled to design an algorithm now prompts the system for code; a literature student asks for an essay plan; a language learner demands instant translation and feedback. The ***automation of reasoning*** thus becomes both a promise and a peril.

This duality revives a long-standing debate in educational philosophy: should technology *extend* cognition or *replace* it? For De La Higuera and Iyer (2025), rapid AI adoption risks reproducing the same techno-centric logic that delayed the meaningful integration of computers in classrooms forty years ago. Without critical frameworks, the risk is not technological failure but cognitive atrophy — a generation of learners proficient in prompting but deficient in thinking.

Methodology

Study Design

To examine how generative AI affects programming skill acquisition, a quasi-experimental design was implemented in the *Introduction to Imperative and Recursive Programming* module at the University of Lorraine. The course aims to develop foundational skills in algorithmic reasoning, problem decomposition, and debugging using the C language.

Two cohorts participated over consecutive academic years (n = 151; ages 18-21). Both followed identical curricula but alternated between ***paper-based exams (CC)*** — unaided, closed-environment tests — and ***computer-based exams (TP)***, with unrestricted access to generative AI and the Internet. This alternation allowed for direct comparison between AI-assisted and unaided performances.

Assessment Instruments

Students completed five assessments per semester:

- ***CC1, CC2, CC3*** - Paper-based, no AI or Internet access.
- ***TP1, TP2*** - Computer-based, open-access with AI allowed.

Each assessment involved three core competencies:

1. ***Problem decomposition***: outlining algorithmic steps;
2. ***Debugging***: identifying and correcting syntax or logic errors;
3. ***Algorithm design***: developing original solutions to novel problems.

Performance was scored on a 20-point scale, and results were analyzed across exam types to evaluate short-term and longitudinal learning effects.

Data Analysis

Statistical analyses included:

- ***Mean and standard deviation*** comparisons between CC and TP scores;
- ***Repeated-measures ANOVA*** to test the impact of exam type;
- ***Pearson correlation*** between AI-assisted and unaided performances;
- ***Chi-square tests*** to examine score distribution shifts across ability levels.

Quantitative findings were complemented by qualitative data from *think-aloud protocols* (n=20), where students verbalized their reasoning during AI use. This enabled exploration of metacognitive engagement and strategic variation in AI-assisted problem-solving.

Limitations

As in most classroom-based studies, several limitations apply:

- ***Sample constraints***: two cohorts from a single institution;
- ***Uncontrolled variables***: prior coding experience and motivation may influence outcomes;
- ***Usage variability***: AI interactions were not logged, limiting analysis of engagement depth;
- ***Estimated longitudinal effects***: performance in advanced courses was inferred from project outcomes.

Despite these constraints, the design enables a rare mixed-method exploration of AI's pedagogical impact, combining statistical rigor with reflective interpretation.

Results

Descriptive Statistics

Across both cohorts (n = 151), students performed significantly better in AI-assisted computer-based exams (TP) than in unaided paper-based exams (CC).

Exam Type	*Mean Score (/20)*	*Standard Deviation*
CC1	9.97	5.58
CC2	8.94	5.50
CC3	9.47	5.62
TP1	12.07	5.25
TP2	13.29	4.71

Table 1: Comparison of mean exam scores across paper-based (CC) and AI-assisted (TP) assessments.

→ Students achieved a ***20-40% increase*** in performance when AI was available.

Inferential Statistics

A ***repeated-measures ANOVA*** confirmed a statistically significant effect of exam type on performance:

- $F(1,150) = 38.45$, $p < 0.001$.
- ***Post-hoc tests***: TP1 > CC1 ($p < 0.01$), TP2 > CC2 ($p < 0.001$).
- No significant differences across the three CC exams (CC1 ≈ CC2 ≈ CC3).

→ ***Interpretation***: The improvement occurs only in AI-assisted contexts, not over time, suggesting short-term performance gains without lasting learning progression.

Correlation Analysis

Variable Pair	*Pearson's r*	*Significance (p)*
TP mean vs. CC mean	0.15	> 0.05
(TP1–TP2) vs. (CC1–CC3)	0.12	> 0.05

→ ***Interpretation***: Weak correlations indicate minimal transfer of skills from AI-assisted to unaided contexts. Students performing well with AI do not necessarily perform well without it.

Score Distribution

Score Category	*CC Exams (Total)*	*TP Exams (Total)*
Low (≤8)	135	60
Medium (9–12)	165	120
High (≥13)	90	150

Chi² (df = 2) ≈ 28.5, p < 0.001

→ The number of high-performing students nearly ***doubled***, and low-performing students halved, suggesting that AI acts as a ***performance equalizer***, flattening score variability.

Qualitative Findings

Think-aloud data revealed three dominant usage profiles among students:

Usage Strategy	*Percentage (n=20)*	*Description*
Direct code generation	60%	Students requested full solutions from AI with minimal reflection.
Explanatory querying	25%	Students asked AI to clarify concepts or debugging logic.
Mixed strategic use	15%	Students alternated between generation and reflection.

→ ***Interpretation***: Most students used AI as a *shortcut*, not a learning companion. Only a minority displayed reflective, metacognitive use of AI-generated information.

Discussion

Performance Gains Without Learning Gains

The data reveal a striking dichotomy between *performance* and *competence*. Students using AI tools achieved higher scores in computer-based exams, but their improvement did not translate to unaided tasks. This pattern suggests ***cognitive outsourcing*** — students offload the reasoning required for algorithm design and debugging to AI systems.

From a pedagogical standpoint, this supports the hypothesis that AI promotes ***surface learning*** (Entwistle, 2009). While it boosts immediate productivity, it does not foster the conceptual frameworks required for sustained skill transfer. The stagnation of CC results across the semester reinforces this conclusion: exposure to AI did not enhance independent reasoning or algorithmic understanding.

These findings echo a broader concern raised by Gonsalves C. (2024) and the APA Task Force (2024): that generative AI improves *outputs*, not *learning processes*. The absence of correlation between TP and CC performances ($r \approx 0.15$) illustrates a rupture between doing and knowing.

The Cognitive Cost of Automation

Using large language models in educational contexts reduces cognitive load but may also diminish cognitive engagement. A recent MIT neuroimaging study (Kosmyna N. et al. 2025) found that students writing essays with LLM assistance exhibited ***lower neural connectivity*** in regions associated with semantic integration and problem solving compared to those writing unaided. This neuroscientific evidence aligns with the observed pedagogical outcomes: ***reduced mental effort correlates with weaker long-term retention***.

This phenomenon mirrors what Risko and Gilbert (2016) describe as "the outsourcing of cognition". In earlier decades, calculators externalized arithmetic; today, AI externalizes reasoning itself. The danger lies not in using AI, but in ***forgetting to think because the machine thinks for us***.

The Collapse of Bloom's Hierarchy

In Bloom's taxonomy, learning is cumulative: knowledge precedes comprehension, which precedes application, and so forth. Generative AI collapses this sequence. Students can now jump from question to "creation" without passing through the intermediate levels. This *shortcut cognition* undermines what Bloom called the "structure of knowledge" — the deliberate ascent through analytical and evaluative thinking.

Bloom's Level	***Traditional Human Process***	***AI-Supported Equivalent***	***Pedagogical Risk***
Remember	Recall facts	AI retrieves instantly	Loss of memorization discipline
Understand	Explain concepts	AI paraphrases	Superficial grasp
Apply	Use in new context	AI generates examples	No transfer learning
Analyze	Identify structures	AI provides analytical breakdown	Passive understanding
Evaluate	Judge solutions	AI scores or critiques	Delegated judgment
Create	Produce new ideas	AI simulates creativity	Homogenization of outputs

Table 2: The compression of Bloom's cognitive hierarchy under generative AI.

As Vivian (2025) observes, this inversion challenges the very foundations of curriculum design: learning becomes less about constructing knowledge and more about managing outputs. The educational task, therefore, is to ***rebuild reflection into the workflow*** — teaching students to justify, critique, and reinterpret AI suggestions rather than reproduce them.

The Illusion of Mastery

A central danger of AI-assisted learning is what psychologists term the ***illusion of explanatory depth*** (Rozenblit & Keil, 2002): the belief that one understands a concept more deeply than one actually does. When AI provides correct solutions, students may overestimate their competence. This false confidence is pedagogically harmful, as it undermines self-assessment accuracy and motivation for further learning.

In qualitative interviews, several students expressed statements such as "*I understood it because the AI explained it well*" — yet performed poorly on unaided tasks. Their experience typifies what Slimi (2023) describes as *synthetic understanding*: comprehension replaced by coherence recognition. The answer "makes sense," so it feels understood — even when no mental model has been built.

From Dependence to Metacognition

Despite these risks, AI can still serve as a ***cognitive amplifier*** when integrated with reflective practice. Metacognitive strategies — such as prompting students to explain how they verified AI responses or to critique their accuracy — can turn dependency into learning. For instance, ***hybrid assessment models*** (Vivian, 2025) that require students to (1) use AI to solve a problem, then (2) replicate or explain it without AI, encourage awareness of both tool and process.

Such hybridization supports what calls *co-agency*: the joint regulation of learning between human and artificial intelligence. Rather than opposing AI and cognition, this approach redefines learning as ***negotiation*** — a dialogue between what the human knows and what the machine suggests.

Ethical and Social Implications

Beyond individual cognition, AI integration raises systemic ethical questions. Plagiarism detection, authorship attribution, and data privacy are pressing issues (Zawacki-Richter et al., 2023). But equally critical is the ***pedagogical ethics*** of dependency. If students increasingly rely on AI to produce, assess, and even think, education risks devolving into *outsourced cognition*.

The long-term societal consequence could be a "generation of editors" — proficient at refining machine output but less capable of generating original thought. As Bloom's model aimed to develop *autonomous thinkers*, AI's automation of reasoning may invert this aspiration, producing ***efficient yet intellectually fragile learners***.

Pedagogical Implications and Future Directions

Rethinking the Role of the Educator

In the age of generative AI, the teacher's role shifts from transmitting knowledge to ***orchestrating cognition***. As LLMs increasingly handle content production, educators must guide students in *evaluating, verifying, and contextualizing* AI outputs. This echoes Vygotsky's (1978) notion of the ***zone of proximal development*** — but with the machine now occupying the role of the "more capable other". The educator's task is to mediate this relationship, ensuring that students remain agents of their own learning.

Practical implications include integrating AI not as an assistant, but as a *topic* of learning. Teaching students to craft effective prompts, question biases, and verify outputs transforms AI from a shortcut into a critical thinking catalyst. Courses that explicitly include "AI literacy" — critical awareness of how generative systems work and where they fail — may help restore cognitive depth in digital learning contexts (Daher, R. 2025).

Designing Hybrid Pedagogical Models

Traditional assessment structures, centered on individual performance, are increasingly incompatible with AI-saturated environments. New models are emerging that blend ***AI-assisted and non-assisted tasks*** within the same learning sequence. For example, a programming course might include:

1. ***AI-augmented creation***: students use AI to generate code or designs.
2. ***Critical replication***: students reproduce and explain outputs without AI.
3. ***Collaborative reflection***: group analysis of AI errors and biases.

Such models foster ***metacognitive awareness*** — students must not only produce results but understand their origin and validity. This approach aligns with the framework of *augmented intelligence* rather than *artificial intelligence* (Annuš N. 2024), viewing the machine as an enhancer of thought rather than its replacement.

Fostering Epistemic Vigilance

AI systems excel at fluency but not necessarily at truth. Their persuasive tone can obscure factual errors or logical gaps (Bender et al., 2021). To counter this, educators must cultivate ***epistemic vigilance*** — the capacity to question and verify claims — as a core learning outcome. Activities emphasizing *argument reconstruction*, *source tracing*, and *bias detection* can transform AI use into epistemic training rather than cognitive surrender.

This pedagogical turn repositions critical thinking as a survival skill in the information age. As Peñalvo (2023) argue, students must learn to operate within an *AI epistemology* — one that privileges reasoning over reproduction, discernment over generation.

Cultivating Creativity Through Constraint

Paradoxically, the abundance of generative capacity may reduce authentic creativity. When AI provides infinite outputs, originality risks dissolving into recombination. Studies in design education (Chakraborty et al., 2024) show that imposing ***creative constraints***

— such as limiting prompts, restricting data access, or requiring manual post-processing — can help preserve divergent thinking.

Educators can thus employ *creative friction* as a learning device. Asking students to critique or extend AI-generated solutions encourages innovation rather than imitation. Creativity becomes the ability to *transcend* machine output, not merely to curate it.

Toward Ethical Co-Creation

Finally, the integration of AI in education raises urgent questions of ethics and authorship. The European Commission's *Ethical Guidelines for AI in Education* (2022) emphasize transparency, accountability, and human oversight. Institutions must develop clear frameworks distinguishing legitimate AI-assisted work from plagiarism, while encouraging reflective documentation of AI use.

Ethical co-creation implies that learners explicitly acknowledge AI's contribution to their work, articulate their own input, and reflect on the interplay between both. This practice not only prevents academic misconduct but also deepens students' understanding of authorship, agency, and responsibility in hybrid human-machine contexts.

Conclusion

Generative AI stands at the frontier of educational transformation — a tool of immense potential and profound ambiguity. The empirical evidence presented here demonstrates that while AI enhances *short-term performance*, it offers limited benefits for *long-term learning*. Students excel with assistance but falter without it, revealing a troubling gap between externalized output and internalized understanding.

At the theoretical level, AI disrupts the foundations of Bloom's taxonomy by collapsing the stages of cognitive development. What was once a gradual climb from remembering to creating becomes an instantaneous leap from prompt to product. In doing so, AI reshapes not only how we learn but what learning itself means. The risk is not ignorance, but ***intellectual complacency*** — a world where knowing is replaced by generating.

Yet the same technology that threatens cognition can also ***amplify it***. When integrated with reflection, constraint, and ethical awareness, generative AI can become a catalyst for deeper learning — a mirror through which students confront their own understanding. The pedagogical challenge is to design environments that foster *co-agency* between human and machine, where the learner remains the author of their reasoning.

As education enters the era of co-intelligence, the goal is no longer to teach students to *think like machines*, but to think ***with*** them — critically, creatively, and consciously. Only then can AI fulfill its promise not as the end of learning, but as its most demanding evolution.

References

1. Anderson LW and Krathwohl DR. (Eds.). "A taxonomy for learning, teaching, and assessing: A revision of Bloom's taxonomy of educational objectives". Longman (2001).
2. Annuš N. "Education in the Age of Artificial Intelligence". TEM Journal 13.1 (2024): 404-413.
3. Bender EM., et al. "On the dangers of stochastic parrots: Can language models be too big?". Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (2021): 610-623.
4. Bloom BS., et al. "Taxonomy of educational objectives: The classification of educational goals". Longmans, Green (1956).
5. Chakraborty P, Ranjan P and Shen L. "Creative constraints in AI-assisted design education". Computers & Education 206 (2024): 105093.
6. Chen A., et al. "The Impact of Generative AI-Powered Code Generation Tools on Software Engineer Hiring: Recruiters' Experiences, Perceptions, and Strategies arxiv.org (2024).
7. Gonsalves C. "Generative AI's Impact on Critical Thinking: Revisiting Bloom's Taxonomy". Sage Journal (2024).
8. Daher R. "Integrating AI literacy into teacher education: a critical perspective paper". Discov Artif Intell 5 (2025): 217.

9. Deci EL and Ryan RM. “Self-determination theory: Basic psychological needs in motivation, development, and wellness”. New York 2017: Guilford Press (2019): 756.
10. De La Higuera C and Iyer S. Book Review: AI for Teachers: An Open Textbook— Artificial Intelligence for and by Teachers International Review of Research in Open and Distributed Learning 26.2 (2025).
11. European Commission. “Ethical guidelines on the use of artificial intelligence (AI) and data in teaching and learning for educators”. Publications Office of the European Union (2022).
12. Peñalvo FJ. “Generative Artificial Intelligence. Open Challenges, Opportunities, and Risks in Higher Education”. In Proceedings 14th International Conference on eLearning 3696 (2023): 4-15.
13. Risko EF and Gilbert SJ. “Cognitive offloading”. Trends Cogn Sci 20.9 (2016): 676-688.
14. Rozenblit L and Keil F. “The misunderstood limits of folk science: An illusion of explanatory depth”. Cognitive Science 26.5 (2002) 521-562.
15. Kosmyna N., et al. “Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task”. arXiv (2025).
16. Shalu., et al. “The Cognitive Cost of AI: How AI Anxiety and Attitudes Influence Decision Fatigue in Daily Technology Use”. Annals of Neurosciences (2025).
17. Slimi Z. “The Impact of Artificial Intelligence on Higher Education: An Empirical Study”. European Journal of Educational Sciences 10.1 (2023).
18. Sparrow B, Liu J and Wegner DM. “Google effects on memory: Cognitive consequences of having information at our fingertips”. Science 333.6043 (2011): 776-778.
19. Vivian R. “From Artificial Intelligence to Real Dummies”. Digital Technologies Research and Applications (2025).
20. Vygotsky LS. “Mind in society: The development of higher psychological processes”. Harvard University Press (1978).
21. Zawacki-Richter O, Kerres M and Bond M. “Artificial intelligence and higher education: Current research trends and future directions”. International Review of Research in Open and Distributed Learning 24.3 (2023): 89-113.