International Journal of Research and Review in Applied Science, Humanities, and Technology
Vol 3, Issue 4 Oct-Dec 2026 ISSN: 3048-975X
https://ijrasht.com/

The Cost of Convenience in Human-AI interaction: Cognitive Offloading and Skill Reconfiguration

Samriddhi Gautam, Department of AIML, Indore Institute of Science & Technology, Indore, Madhya Pradesh, India samriddhi.gautamaiml2024@indoreinstitute.com

Samarth Dubey, Department of AIML, Indore Institute of Science & Technology, Indore, Madhya Pradesh, India samarth.dubeyaiml2024@indoreinstitute.com

Praveena Joshi, Department of AIML, Indore Institute of Science & Technology, Indore, Madhya Pradesh, India praveena.joshi@indoreinstitute.com

Shreyas Pagare, Department of AIML, Indore Institute of Science & Technology, Indore, Madhya Pradesh, India shreyas.pagare@indoreinstitute.com

ABSTRACT: Generative AI, and more precisely large language models, have nowadays rapidly become useful in education, professional work, and everyday tasks. While these systems offer notable benefits to make work efficient and improved accessibility, the recent evidence suggests that their usage has cognitive costs. This paper presents a structured review of 16 reviewed studies investigating the impact of Gen AI usage on cognitive offloading, metacognitive engagement, automation bias, and human skill. By studying experimental, survey-based, qualitative, and theoretical research, this review synthesizes several findings across several fields to identify recurring mechanisms through which GenAI usage reshapes cognition. Our analysis reveals a consistent pattern, where GenAI often improves short-term task performance, but it frequently minimizes effortful cognitive processing, diminishes self-regulation, and changes skill developments. Rather than constant de-skilling, this paper points out reconfiguration of skills, favoring efficiency-oriented competencies such as prompting and verification over generative reasoning and critical reflection. Our study further highlights these outcomes are not unavoidable but are strongly mediated by interaction design, educational practices, and socio-technical environments. By showing these effects as the "cost of convenience," this paper contributes an integrative perspective on the cognitive effects of Gen AI and highlighting the ways to make AI design, education, and governance cognitively sustainable.

Keywords - Generative AI, AI dependency, Skill Degradation, Cognitive Offloading, Cognitive Skills.

Introduction

The pace at which Generative Artificial Intelligence (GenAI) has been developed from an experimental stage to the most used cognitive assistant of our lives. Gen AI nowadays is becoming a common component in our careers, educational settings and in our everyday decisions. The speed at which users can create text, code, explanations, and recommendations, means we have got to see significant improvements in efficiency, accessibility and productivity through the usage of AI chatbots and systems day-by-day. However, significant concerns related to cognitive effects which are associated with assigning our main intellectual responsibilities to AI technologies, have begun to take place.

One of the main characteristics of Gen-AI is its convenience: things that previously required sustained reasoning, memory retrieval, or creative synthesis which can now be done by simple prompts. While this convenience shrinks the borders for completing a task, it also raises fundamental questions such as: What happens to human cognitive skills when effortful thinking is systematically bypassed? This question has been the focus of recent discussions on cognitive offloading, over-reliance on AI, and the potential degradation of critical thinking, metacognition, and the problem-solving ability.

Earlier research on cognitive offloading has shown that various external tools can support as well as undermine cognition, depending upon how they are used. Unlike traditional tools, Gen-AI not only stores or retrieves information—it actively performs functions like reasoning. Qualitative changes like this increase concerns that if frequently relied on Gen-AI. It can reduce opportunities for cognitive practices, foster automation biases, and can reshape learning and work habits in ways that are difficult to turn back. Some previous studies report improved short-term performance under assistance of AI, but it simultaneously weakens learning, reduces cognitive effort, and diminishes self-regulation [1], [4], [15].

Despite a rapidly expanding work, existing research remains fragmented across disciplines, including the interactions between humans and computers in fields of education, psychology, and socio-technical studies. Many studies focus on isolated outcomes, or specific user groups, or single task domains, and it gets hard to draw integrated conclusions about the wider cognitive implications of the usage of Gen-AI. However, this literature combines the gains in efficiency with cognitive benefits, overlooking the possibility that this 'convenience' of using Gen-AI may carry hidden long-term cognitive costs.

Our review recognizes these gaps by proper examination of evidence leading towards the cognitive offloading and human skill degradation by the usage of AI, as the cost of convenience. Instead of evaluating Gen-AI systems solely over performance, our review focuses on underlying cognitive mechanisms and skill trajectories. By combining the empirical, theoretical, and structural perspectives, this paper provides a comprehensive account of how usage of Gen-AI influences an individual's cognition at educational, organizational, and societal levels.

Contributions of This Review

Contributions made by this paper:

1. **Conceptual synthesis:** It organizes existing research around the key cognitive structures—cognitive offloading, metacognitive laziness, automation bias, and skill reconfiguration—providing a combined explanatory framework.
2. **Comparative analysis:** It systematically compares methods of research and AI interaction techniques to clarify why findings converge across diverse methodologies.
3. **Critical perspective:** It reframes the risks of Gen AI not as unavoidable consequences of usage of AI, but as outcomes shaped by interaction design, institutional incentives, and socio-technical environments.
4. **Design and policy relevance:** It identifies ways to mitigate cognitive costs through cognitively sustainable AI design, educational practices, and governance strategies.

This paper is organized as follows: Section II describes the methodology and literature identification and screening process. Section III synthesizes the literature according to core cognitive mechanisms. Section IV shows a comparative analysis of methods of research and Human-AI interaction techniques. Section V discusses cross-study implications, limitations, and future research directions. At last, Section VI concludes by summarizing key insights and in boxing the broader significance of understanding the cognitive cost of usage of Gen-AI.

Methodology

The main objective of our review is to analyze how using Generative Artificial Intelligence (GenAI) affects human cognition and skill development. To get to this, a proper literature identification and screening process was followed to ensure relevancy, transparency, and consistent analytics.

Literature Identification and Screening

Only relevant studies were identified through searches across trusted and authentic academic platforms, like Scopus, Web of Science, IEEE Xplore, ACM Digital Library, ScienceDirect, and Google Scholar. Search queries were combined with GenAI-related terms, cognition- and skill-oriented bases to restrict results for relevant outcomes.

Only studies published between 2019 and early 2025 were considered in this review, keeping the period in mind when large language models and related GenAI tools became widely accessible and began influencing routine cognitive activities. The review is limited to English-language journal articles and conference papers. Preprints were examined very selectively and kept only when no peer-reviewed alternative existed or could not be found.

Screening was conducted in sequential stages. Records were first sorted for duplicates and topical relevance using titles and abstracts. Then the complete texts were assessed to verify the clear engagement with at least one of the following themes: cognitive efforts, cognitive offloading, metacognitive regulation, skill degradation, skill reconfiguration, or over-reliance on AI. Studies focusing exclusively on non-generative automation, robotics, or domain-specific AI without cognitive analysis were not included.

This process resulted in a final corpus of 16 papers, representing a mix of experimental research, surveys, qualitative investigations, theoretical analysis, and neuroscientific base approaches. Rather than covering more things, our review prioritizes the analytical depth and conceptual coherence of the topic,by acknowledging the rapidly evolving tools and unpublished industry reports that fall outside the present scope. The applied selection constraints are summarized in Table I.

TABLE I - Search-Criteria for Selection of Literature

Criteria	Description
Keywords Used	"Generative AI", "ChatGPT", "Cognitive Offloading", "Critical Thinking", "Metacognition", "Automation Bias", "AI Over-Reliance", "Skill Degradation", "Deskilling", "Human–AI Collaboration"
Publication Type	Peer-reviewed journal articles and conference proceedings
Databases / Indexes	Scopus, Web of Science, IEEE Xplore, ACM Digital Library, ScienceDirect, Google Scholar
Time Span	2019–2025
Language	English
Total Papers Identified	~1,064
Inclusion Criteria	Studies examining Generative AI usage and its impact on cognition, learning, critical thinking, metacognition, or skill development; empirical, theoretical, or

	systematic review studies
Exclusion Criteria	Opinion pieces, editorials, blogs; non-GenAI automation studies; robotics-only research; medical AI without cognitive focus; non-English publications
Final Papers Selected	16

Conceptual Synthesis Framework

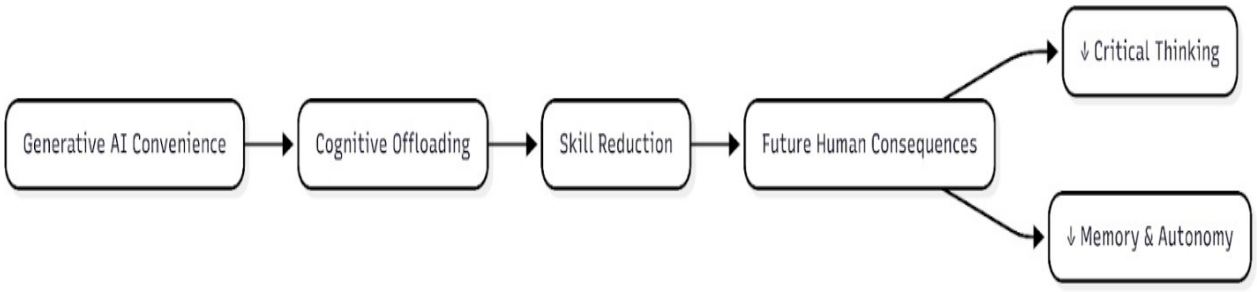


Figure shows the framework we used to make sense of all the research reviewed. Basically, it maps out a chain reaction: when people use GenAI because it's easy and convenient, they start offloading their thinking to the AI. Over time, this weakens their ability to monitor and regulate their own thinking, and they become more likely to just trust whatever the AI says without questioning it. These patterns then shape how people's skills change in the long run—sometimes skills deteriorate, other times they just shift into something different, depending on the situation and how the AI tools are designed. I've used this figure as a roadmap for organizing the literature review that follows.

Literature Review

We've brought together existing research on Generative AI by focusing on the cognitive mechanisms that drive skill dependency and degradation. Rather than reviewing each study in isolation, We identified four recurring themes that consistently appeared across the literature: cognitive offloading, metacognitive laziness, automation bias, and the distinction between skill degradation and skill reconfiguration. Together, these themes help us understand how the convenience of generative AI fundamentally changes the way we think and learn.

A. Cognitive Offloading in Generative AI Usage

A considerable body of factual evidence reveals that Gen AI systems support cognitive offloading, which can be defined as the imposition of mental tasks such as remembering, reasoning, thinking, problem-solving to external tools [1], [8], [9], [13], [15]. Across experimental, survey-based, and modeling studies, increased dependency on AI generated output is directly associated with less involvement in tasks that include mental activity.

Research trials conducted under regulated conditions demonstrated that when users were given automatic AI-generated solutions, they would be less inclined to carry out independent analytical or problem-solving work [1],[4]. Although the effectiveness of task completion (i.e., speed of completing tasks) and short-term performance should improve while utilizing AI solutions; however, these benefits will be coupled with a significantly lower ability to transfer skills to other situations and will experience reduced overall performance when AI assistance ceases, indicating they created an illusion of competence [1].

In support of this conclusion, larger population studies of survey data indicated that extensive use of AI products is associated with a high-level of Cognitive Offloading and lower levels of Critical Thinking [8],[13],[15]. Additionally, mediation analyses suggest that Cognitive Offloading somewhat accounts for the negative consequences of using AI-products which often have to do with Critical Thinking [8].

Neuroscience and modeling research has additional evidence supporting the earlier findings: AI assisted task environments produce a decrease in measures that assess Executive Cognitive Engagement as well a decrease in Memory Retention [5],[9]. Collectively, these studies may not be consistent in their levels of Ecological Validity; however, they collectively demonstrate that using Generative AI technology results in Cognitive Offloading, thus removing the need for cognitive processing via Internal Process Development, allowing for External Delegation of these processes.

B. Metacognitive Laziness and Reduced Self-Regulation

Beyond task execution, Generative AI use has been shown to erode metacognitive regulation, a phenomenon described as metacognitive laziness [2], [4], [10]. This refers to reduced participation in planning, monitoring, and

evaluating one's own cognitive processes.
Experimental studies examining learning processes illustrate that users supported by Generative AI exhibit fewer metacognitive transitions, particularly in evaluation and reflection phases [4]. Users who receive AI assistance tend to revise faster as they self-question or analyze and detect errors lesser than the ones who receive human responses. Research which is conceptual and synthesis-oriented further suggests that Gen-AI can help learners develop a solution-first approach, where the learners will primarily focus on getting the correct answers rather than understanding underlying principles [2], [10]. As a result, people do less self-directed learning and lose the chance to grapple with difficult problems—the kind of struggle that actually builds skills.
More importantly, multiple studies say that metacognitive laziness comes not from lack of motivation but from user-interaction designs that minimizes cognitive friction and make the need for reasoning less secure [2], [4].

C. Automation Bias and Over-Reliance on AI Outputs

The third theme focuses on automation bias, where users blindly trust the AI-generated outputs and fail to critically evaluate their correctness [6], [11], [15]. Generative AI systems are particularly exposed to this effect due to their fluent, confident, and human-like responses.
Some review-based studies constantly report that users are more likely to accept AI suggestions even when they are incomplete, biased, or incorrect [6], [11]. This uncritical acceptance is increased when tasks are perceived as low-stakes or when AI outputs align with users' past beliefs.
Recent survey studies among knowledge workers highlight the role of confidence calibration in automation bias. Higher confidence in AI systems is associated with decrease in critical thinking, whereas confidence in one's own abilities predicts greater cognitive engagement [15]. Over time, these results normalize verification-only workflows, shifting users from active reasoning to passive oversight.

D. Skill Degradation versus Skill Reconfiguration

Regardless of whether most of this review is focused on skill degradation, recent work highlights a more refined outcome: skill reconfiguration rather than uniform loss [7], [12], [14], [16]. Foundational skills such as critical reflection, deep comprehension, and creative synthesis are frequently under-practiced when AI performs core cognitive functions [7], [12], [16].
Simultaneously, newer efficiency-oriented skills come out, such as prompt writing, output evaluation, and rapid integration of
AI-generated content [14], [16]. Qualitative evidence suggests that users are often aware of this trade-off but accept it due to time and performance pressure [16].
Structured comprehensive analysis argues that these shifts are backed up by capacity-hostile environments, where AI systems arrange the user as responder instead of the initiator who performs the cognitive activity [3], [7], [12]. Over time, this may worsen the inequalities, as individuals and institutions with stronger cognitive foundations are better positioned to benefit from AI without depending on it [3].

Table II - Comparison Table

Paper (Year)	Problem Addressed	Dataset Used	Features / Model Type	Evaluation Metrics	Key Results	Limitations
Kazemitabaar et al. (2025)	Whether direct AI code generation reduces learning via cognitive offloading	2 user studies (N=124 CS undergraduates)	GPT-4 code generation; engagement-inducing interaction designs	Coding scores, transfer tasks, NASA-TLX, time	Direct AI use caused illusion of learning; engagement techniques improved skill retention	Short-term, CS-only tasks, student population
Giannakos et al. (2025)	Risks of over-automation in GenAI-supported education	Expert synthesis, prior studies	LLMs (ChatGPT, Copilot); hybrid intelligence framing	Conceptual analysis	GenAI convenience can undermine SRL and deep learning	No empirical validation
Cazzaniga et al. (2025)	Unequal skill and productivity effects of AI	IMF, ILO, OECD datasets (174	Macroeconomic simulation (TFP shocks)	GDP, TFP growth projections	AI benefits skilled economies; weak preparedness leads to dependency	No micro-level cognition analysis

	adoption	countries)				
Fan et al. (2025)	Whether ChatGPT induces metacognitive laziness	RCT (N=117 students)	ChatGPT vs human expert vs checklist tools	SRL process mining, essay scores, IMI	AI improved performance but reduced metacognitive engagement	Short-term, writing-only task
Chen et al. (2025)	Does GenAI reduce cognitive effort or reallocate it?	Planned RCT (N=160 students)	GPT-4; eye-tracking & fNIRS	Pupillometry, fNIRS, essay scores	Results pending; framework targets causal cognitive effects	Protocol only; no results
Nah et al. (2023)	Over-reliance risks in AI–human collaboration	Conceptual + case synthesis	ChatGPT, GenAI systems	Conceptual dimensions	AI convenience increases automation bias and dependency	Non-empirical
Sætra (2023)	Societal and cognitive risks of GenAI	Philosophical analysis	LLMs, generative systems	Normative reasoning	Warns of cognitive atrophy and autonomy loss	No empirical testing
Gerlich (2025)	Does AI reduce critical thinking via offloading?	Survey (N=666) + interviews	AI tools broadly (assistants, decision aids)	HCTA, regression, Random Forest	Cognitive offloading mediates critical thinking decline	Cross-sectional; self-report
Narayan (2025)	Neural and behavioral cost of AI convenience	Simulated EEG dataset (N=120)	Conceptual GenAI use conditions	EEG activation, recall, CT scores	Higher AI reliance → lower neural engagement & recall	Simulated data; limited validity
Singh et al. (2025)	How to protect cognition in GenAI era	Integrative literature review	LLMs across domains	Synthesis of prior metrics	GenAI reduces reflective thinking without safeguards	Theoretical
Zhai et al. (2024)	Over-reliance effects of AI dialogue systems	Systematic review (14 studies)	AI chatbots, GenAI tools	Prior CT & reasoning measures	Consistent evidence of critical thinking decline	Limited domains; no longitudinal data
Ferdman (2025)	Why AI deskilling is structural, not individual	Philosophical analysis	LLMs, APAs, recommender systems	Conceptual evaluation	AI creates capacity-hostile environments	Non-empirical
Ejaz et al. (2025)	Cognitive load reduction vs critical thinking	Survey (N=350 students)	ChatGPT, AI tools	Regression, CLT scales	Lower cognitive load predicts lower CT	Self-reported CT
Kabashkin (2025)	Cognitive Atrophy Paradox of AI use	Simulations + case illustration	LLMs; Cognitive Sustainability Index	CSI, growth/atrophy models	Convenience shifts cognition toward dependency	Needs empirical validation
Lee et al. (2025)	GenAI impact on critical thinking at work	Survey (N=319 knowledge workers; 936 tasks)	ChatGPT, Copilot, Gemini	Mixed-effects regression, Bloom taxonomy	GenAI reduces cognitive effort; confidence in AI reduces CT	Self-report; no longitudinal data
Passlack et al. (2025)	Deskilling vs reskilling in academic research	Interviews (N=12 students)	ChatGPT in research projects	Gioia qualitative coding	Efficiency-driven GenAI use reduces creativity	Small sample; qualitative

Comparative Analysis

While the last section was around results and shared cognitive mechanisms, this section compares the reviewed studies along two critical dimensions:

(1) Research methodologies used to study cognitive offloading.
(2) AI interaction techniques and usage styles that shape cognitive outcomes.

These comparisons help in explaining the results across different research domains and also the sources of inconsistencies.

A. Comparison of Research Methods

The reviewed literature acknowledges a wide range of methodological approaches to study how Generative AI impacts cognition. Table III compares these approaches in terms of evidence strength, scalability, and limitations. Rather than prioritizing a single method, this comparison table highlights how all the evidence merges to strengthen the claim of cost of convenience.

Table III: Comparison of Research Methods in Studies of Cognitive Offloading

Methodological Approach	Primary Data / Metrics	Strengths	Limitations	Representative Studies
Controlled Experiments	Task performance, learning transfer, process traces	Causal inference; fine-grained cognitive analysis	Short-term; lab settings; limited generalizability	[1], [4]
Large-Scale Surveys	Self-reported AI use, critical thinking scales, regression models	High scalability; population-level trends	Self-report bias; cross-sectional	[8], [13], [15]
Qualitative Studies	Interviews, thematic coding, lived experiences	Rich contextual insight; mechanism discovery	Small samples; no causal claims	[16]
Neuroscience/ Modeling	EEG proxies, eye-tracking, simulations	Direct cognitive effort indicators	Limited ecological validity; simulated data	[5], [9], [14]
Systematic Reviews	Synthesized empirical findings	Consolidated evidence; pattern identification	Dependent on prior study quality	[11]
Conceptual/ Structural Analyses	Normative reasoning, socio-technical theory	Explains persistence and scale of effects	No direct measurement	[2], [6], [7], [12]

Key insight: Although individual methods have limitations, the consistency of results across experimental, survey-based, qualitative, and theoretical studies provides strong cumulative evidence that Gen-AI usage promotes cognitive offloading and skill reduction.

B. Comparison of AI Interaction Techniques and Usage Styles

After methodologies, studies differ on how Generative AI is used for cognitive tasks. Table IV compares AI interaction techniques according to their cognitive impact and the outcomes are determined by specific design choices, not just by the presence of AI.

Table IV - Comparison of AI Interaction Techniques and Cognitive Effects

AI Interaction Technique	Description	Cognitive Impact	Observed Outcomes	Representative Studies
Direct Solution Generation	AI provides complete answers or outputs	High cognitive offloading	Faster task completion; weaker learning transfer; illusion of competence	[1], [8], [13], [15]

Guided Scaffolding	AI prompts reasoning steps or partial hints	Moderate engagement	Improved understanding with controlled effort	[1], [10]
Verification-Orien ted Use	User checks or edits AI outputs	Reduced generative reasoning	Shift from production to oversight; automation bias	[6], [11], [15]
Human-in-the-Loop Support	AI augments but does not replace reasoning	Sustained metacognition	Better balance between efficiency and learning	[2], [4]
Unrestricted AI Assistance	Free-form AI use without constraints	Variable, often negative	Dependency patterns; metacognitive laziness	[4], [16]

Key insight: AI interaction techniques that maximize convenience constantly make the cognitive reduction effects stronger, On the other hand designs that introduce cognitive friction or scaffolding mitigates the skill degradation without losing AI convenience.

Result and Discussion

This paper inspected the cognitive cost of convenience in Gen AI usage which gradually leads to cognitive offloading and skill change. Strong evidence from the experimental, survey based, qualitative and theoretical studies show how recurring trends affect how AI is regulated and taught.

A. The Cost of Convenience: Cross-Study Synthesis

Across domains and methodologies, the literature gathers on a central finding that Generative AI systems reduce effortful cognitive engagement than traditional AI when used for convenience, they even help users to improve short-term efficiency and task performance [1], [4], [8], [13], [15]. This trade-off is a fundamental characteristic rather than accidental. Since AI generated solutions take minimal efforts from the user, they cannot develop and flourish foundational skills.

Multiple reinforcing mechanisms develop cognitive offloading and it eventually results in weakened metacognitive and self-evaluation skills [2],[4], also lack of practice leads to reduced foundational skills [1],[8], automation bias comes in the frame which leads to blindly accept AI outputs [6],[11],[15]. Ultimately, these systems create a loophole where users fall for ease and speed but gradually lose their skillset.

While studies of artificial intelligence generally indicate improved performance outcomes for those who receive AI assistance, there are significant negatives associated with AI use as well. Due to this, while AI may provide a short-term advantage to the user with respect to performance, it also has numerous long-term disadvantages which may disrupt skill building over time.

B. Skill Degradation as Reconfiguration Rather Than Immediate Loss

Research suggests that rather than uniform deskilling from GenAI use, there is a reconfiguration of skills related to AI use [7], [14], [16]. Users develop efficiency-oriented skills (e.g., prompting, verifying, and integrating outputs) as AI takes over generative cognitive tasks. While these skills are important, they often replace deeper skills (e.g., problem formulation, creative synthesis, and independent judgement).

Qualitative and structural analyses indicate that the context and its incentives shape reconfigured skills. Many users perceive that they have reduced their deep engagement with AI but have not stopped using it because of their time pressures, expectations from their assessments, and the norms of their organisations which place a higher priority on efficiency than understanding [16]. Thus, skill degradation is explained by an environment that creates barriers to developing or maintaining high levels of cognitive processing ability, rather than an individual's lack of effort [12].

From a macro-structural perspective, this reconfiguration of skills creates a situation that produces greater inequality since individuals and organisations that have stronger cognitive foundations are at a greater advantage when using AI than those who do not [3]. Therefore, the distribution of the costs associated with convenience varies by organisation.

C. Implications for Design, Education, and Policy

The evidence suggests that the cognitive impact of generative AI is largely driven by its design features rather than emerging naturally. It has been demonstrated that interaction design strategies like cognitive scaffolding, delayed solution exposure, and reflective prompting reduce cognitive offloading without compromising effectiveness [1], [4], and [10]. Therefore, to create AI systems that are cognitively sustainable and assess their impact on the advancement of human skills, we argue for a reevaluation of design practices that are mainly concerned with

convenience-maximization. In educational settings, the literature indicates a disconnect between many current assessment practices and emerging AI-based learning environments [2],[11]. When assessments focus on the product (i.e. the final answer) rather than the process by which students arrive at a product, it is reasonable to expect increased dependence upon AI, while at the same time incurring a cognitive toll. Consequently, it is essential that process-oriented assessment, metacognitive instruction, and AI literacy be employed to prevent long-term erosion of skills.

To address deskilling from a governance perspective, it should not be expected that individual users bear full responsibility for this issue
[7],[12]. Institutions and platform providers must evaluate whether AI systems create an environment conducive to either capacity development or capacity reduction; thus, the focus of the policy debate should be on the responsible regulation of convenience rather than the outright prohibition of convenience.

D. Limitations and Future Directions

While growing evidence is converging to support the relationships between cognitive offloading practices and associated cognitive changes, limitations still exist. For example, most of the empirical literature available at present primarily consists of either short-term, cross-sectional or laboratory-based studies, thus limiting the generalizability of findings regarding future cognitive change possibilities. Additionally, most studies use subjective self-reported cognition metrics, which may not always accurately reflect true cognitive changes occurring within participants.

Incorporating some recent emerging research that utilizes both physiological and modeling techniques, however, there are still few established longitudinal measures that provide evidence of cognitive degradation as it relates to offloading practices. As we move toward future research are necessary to understand these relationships further...also will require the utilization of sufficient numbers of mixed-methods studies using longitudinally collected data regard the behavioral, cognitive and neural correlates of cognitive offloading...and will need to include ample large-scale evaluations of the design interventions in real-world settings. Figure II summarizes the synergism between the research findings presented.

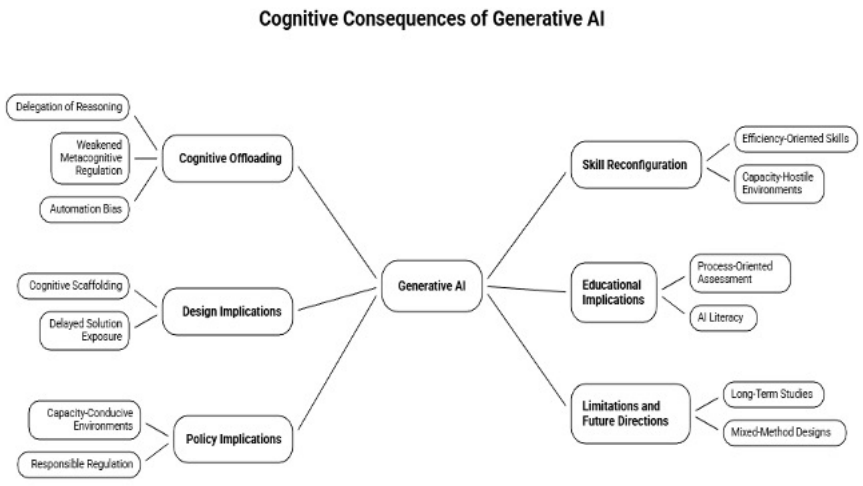


Figure II, Conceptual Framework for the Cognitive Impact from Gen AI - Shows the summary of the literature on generative artificial intelligence as well as the many cognitive impacts due to using generative artificial intelligence, how cognitive offloading (due to delegating reasoning, the weakening of one's metacognitive ability, and creating an automation-based bias) and how cognitive skill reconfiguration translates into cognitive consequences. The findings of this framework provide evidence that the design (cognitive scaffolds), educational system (assessing cognitive processes), and legislatures will need to adapt the current models to account for the cognitive impacts created by the convenience associated with using generative artificial intelligence and that the field has significant limitations regarding the need for longitudinal, multi-method studies of generative artificial intelligence's cognitive consequences.

Conclusion

This article reviews and synthesises new research from 16 high-quality peer-reviewed sources investigating the cognitive impacts of Generative AI technology resulting from the increased convenience associated with using this technology. Generally, across experimental studies, surveys, qualitative studies and studies informed by neuroscience and theory, there were consistent findings that Generative AI technologies can improve efficiency

and task effectiveness over the short term while also resulting in decreased cognitive engagement. However, in many cases, improved efficiency directly correlated with an erosion of deeper cognitive skills.

The most important finding from this review is that the cognitive implications of using Generative AI should not be viewed as immediate or uniform deskilling of an individual but rather through the lens of skill reconfiguration. As generative cognitive functions are assumed by AI technologies, users will progressively develop new capabilities (i.e. focus on efficiency) when prompting (causing the AI to generate something); verifying (checking for accuracy); and integrating (developing a final product using AI-generated content). These new capabilities come at a cost, as users are sacrificing key cognitive competencies that require deeper engagement than those used to complete tasks with AI assistance. The cognitive costs associated with using Generative AI are worsened by users' increased reliance on cognitive offloading, metacognitive disengagement and automation bias, creating a never-ending loop in which reliance on Generative AI continues to increase. However, it is critical to note that these outcomes are not inevitable. The degree of cognitive impact associated with using Generative AI is dependent upon multiple factors, including the design of interactions (how the user interfaces with the technology), practices used to assess educational quality of interaction and the wider social and technical context surrounding Generative AI.

This review provides an integrative framework for understanding how Human–AI interaction is affected by the "cost of convenience." Additionally, it highlights the importance of evaluating Generative AI based on its long-term effects on human cognition rather than its fast and rapid results. As Generative AI continues to gain acceptance in both education and the work environment, it should boost and increase the human potential and skills, not diminish them. Therefore, future research should focus on long-term and real-world studies and scalable design interface, which aim to create products that combine efficiency and skills with the continued support of critical human cognition skills.

REFERENCES

[1] S. Kazemitabaar, J. Wang, A. Head, and B. Hartmann, "Exploring the Design Space of Cognitive Engagement Techniques with AI-Generated Code for Enhanced Learning," in Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI), 2025, pp. xxx–xxx, doi: 10.1145/3708359.3712104.

[2] M. Giannakos, P. Dillenbourg, M. Brundin, K. Sharma, and D. Gašević, "The Promise and Challenges of Generative AI in Education," Behaviour & Information Technology, vol. 44, no. 11, pp. 2518–2544, 2025, doi: 10.1080/0144929X.2024.2394886.

[3] F. Cazzaniga, A. C. Schoder, and J. L. Valencia, "The Global Impact of AI: Mind the Gap," IMF Working Paper, no. 25/076, International Monetary Fund, 2025.

[4] X. Fan, Y. Li, and J. Wang, "Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance," British Journal of Educational Technology, early access, 2025, doi: 10.1111/bjet.13544.

[5] Y. Chen, M. Bießmann, and K. Söllner, "Effects of Generative Artificial Intelligence on Cognitive Effort and Task Performance: Study Protocol for a Randomized Controlled Experiment among College Students," Trials, vol. 26, no. 244, 2025, doi: 10.1186/s13063-025-08950-3.

[6] F. F.-H. Nah, X. Zheng, and L. Cai, "Generative AI and ChatGPT: Applications, Challenges, and AI–Human Collaboration," Journal of Information Technology Case and Application Research, vol. 25, no. 3, pp. 277–304, 2023, doi: 10.1080/15228053.2023.2233814.

[7] H. S. Sætra, "Generative AI: Here to Stay, but for Good?" Technology in Society, vol. 75, p. 102372, 2023, doi: 10.1016/j.techsoc.2023.102372.

[8] M. Gerlich, "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking," Societies, vol. 15, no. 1, p. 6, 2025, doi: 10.3390/soc15010006.

[9] N. Narayan, "AI, Cognition, and the Cost of Convenience," International Journal of Technology and Systems, vol. 7, no. 6, pp. 18–34, 2025.

[10] R. Singh, A. Kumar, and P. Mehta, "Protecting Human Cognition in the Age of AI," arXiv preprint arXiv:2502.12447, 2025.

[11] C. Zhai, S. Wibowo, and L. D. Li, "The Effects of Over-Reliance on AI Dialogue Systems on Students' Cognitive Abilities: A Systematic Review," Smart Learning Environments, vol. 11, no. 28, 2024, doi: 10.1186/s40561-024-00316-7.

[12] A. Ferdman, "AI Deskilling Is a Structural Problem," AI & Society, early access, 2025, doi: 10.1007/s00146-025-02686-z.

[13] A. Ejaz, M. Farhan, and F. E. Alessi Longa, "AI and Cognitive Load: How Reliance on AI Tools (ChatGPT, Etc.) Affects Critical Thinking," Research Journal of Psychology, vol. 3, no. 4, pp. 1–10, 2025.

[14] I. Kabashkin, "Cognitive Atrophy Paradox of AI–Human Interaction: From Cognitive Growth and Atrophy to Balance," Information, vol. 16, no. 11,
p. 1009, 2025, doi: 10.3390/information16111009.

[15] H.-P. Lee, A. Sarkar, L. Tankelevitch, I. Drosos, S. Rintel, R. Banks, and N. Wilson, "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects from a Survey of Knowledge Workers," in Proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI), 2025, pp. Xxx–xxx.

[16] N. Passlack, K.-H. Gerholz, and P. Schlottmann, "Reskilling Me Softly: Perceived Changes in Students' Skilling When Using GenAI in Academic Research Projects," in Proceedings of the 58th Hawaii International Conference on System Sciences (HICSS), 2025.