



Disappearing Minds in the Age of Artificial Intelligence

Matjaz Perc
Mahmut Özer

Artificial intelligence has now become an integral part of daily life for people of all ages. In particular, generative AI tools that utilize large language models (LLMs), such as ChatGPT, are increasingly being used by students and professionals from various fields in their everyday tasks (Ilikhan et al., 2025; Özer, 2024a; 2024b; Perc et al., 2019). With capabilities such as text generation, translation between languages, graphic creation, summarization, and providing advice on encountered problems, these applications offer individuals a highly convenient option. Especially in the field of education—from primary to higher education—many students now use these tools as a significant support in their learning processes (Özer, 2024a; Suna & Özer, 2025; Tanberkan et al., 2024).

Education systems are among those where the transformative impact of artificial intelligence (AI) will be felt most strongly. These systems bear two major responsibilities in response to this transformation. The first is to train individuals with the skills demanded by labor markets. As AI rapidly transforms industries and thus labor markets, the challenge of developing a workforce with skills suited to new conditions remains a pressing issue (Özer & Perc, 2024; Özer et al., 2024a). For example, China has already begun implementing proactive initiatives aimed at equipping students with AI-related skills at the primary, middle, and high school levels. Similar transformations are also taking place in higher education.

The second responsibility relates to the direct impact of these technologies on education systems. In particular, with generative artificial intelligence technologies, it

 Faculty of Natural Sciences and Mathematics, University of Maribor; Community Healthcare Center Dr. Adolf Drolc Maribor; University College, Korea University; Department of Physics, Kyung Hee University; matjaz.perc@gmail.com
National Education, Culture, Youth and Sports Commission, Turkish Grand National Assembly, Türkiye, mahmutoz2002@yahoo.com

 <https://orcid.org/0000-0002-3087-541X>
<https://orcid.org/0000-0001-8722-8670>

 DOI: 10.12658/E0004
insan & toplum, 2025; 15(3): 1-9.
insanvetoplum.org

has become much easier to enrich educational environments and create personalized learning opportunities (Özer, 2024a). On the other hand, significant opportunities are emerging in areas such as foreign language instruction, visual arts, music, and arts education. Moreover, these technologies can provide substantial support in tasks such as generating text, developing project proposals, conducting literature reviews, and writing academic papers (Özer, 2024b; Suna & Özer, 2025).

Of course, like all general-purpose technological transformations, these tools come with not only benefits but also significant risks and challenges. While discussions around these risks often focus on issues such as data security, the amplification of inequality through biased outputs, lack of transparency in decision-making, outdated or incorrect information (so-called “AI hallucinations”), and misinformation generation (Ilikhan et al., 2024; Özer et al., 2024b; Özer, 2024c), more recently, studies have begun to emerge addressing the cognitive costs of these applications. Therefore, this Letter to the Editor evaluates three recent studies that explore the cognitive effects and emerging costs of large language models (LLMs) (Ahmad *et al.*, 2023; Kosmyrna et al., 2025; Stadler et al., 2024). While each study approaches the topic from a different angle, their findings are mutually reinforcing. In light of these findings, the letter proposes new policy recommendations regarding the use of artificial intelligence within education systems.

In this context, the first study examines the effects of artificial intelligence on personal security and privacy concerns, loss of decision-making ability, and increased laziness among a group of higher education students from China and Pakistan (Ahmad *et al.*, 2023). The study’s findings show that, in the student sample used, AI applications have a significant impact on students’ loss of decision-making ability and increased laziness. According to the results, 68.9% of the students’ laziness, 68.6% of their personal privacy and security issues, and 27.7% of their loss in decision-making ability are attributed to the effects of AI applications.

Security and privacy concerns are quite common not only in education but in all fields where artificial intelligence is applied. Moreover, the frequent appearance of news in the media about security vulnerabilities and privacy breaches continuously raises the level of concern. The most important solution individuals can pursue to address this anxiety is to continuously improve their literacy regarding these technologies (Özer et al., 2024b). A lack of knowledge in this area will make data leaks much more likely.

The findings of the study indicate that as the functions of AI applications increase and their use in education becomes more widespread, students’ trust in these applications in their academic work also grows, which in turn leads to increased

laziness (Ahmad et al., 2023). As students delegate more of their academic tasks to these applications, their reflex to check and verify the generated content weakens. Consequently, students increasingly use less of their analytical thinking skills and cognitive abilities in the production of content. Excessive trust in AI-generated content gradually fosters a dependency on these applications. Ultimately, this dependency leads to laziness among students. In fact, a similar risk also applies to teachers and academic staff. On the other hand, the findings of the study indicate that as the use of AI applications in education becomes more widespread in this way, students' ability to make decisions weakens. This is because instead of making decisions themselves, students increasingly adopt the habit of substituting AI-generated decisions for their own. As a result, their ability to cope with the tensions involved in the decision-making process also diminishes. In the long term, this vulnerability may indirectly have a negative impact on students' resilience in the face of challenges.

Beyond these effects, as partially mentioned above, there are two major risks associated with AI-generated content. The first concerns the risk that AI applications may produce biases based on religion, culture, gender, socioeconomic status, race, and similar factors (Özer et al., 2024a; Suna & Özer, 2025). When students place excessive trust in such content, they are unlikely to verify it, and as a result, perspectives that reinforce existing inequalities may become widespread through the education system. The second risk relates to the phenomenon of hallucination or confabulation by AI applications (Özer, 2024c). When such behavior occurs, AI systems generate false content that appears to be accurate. If this content is not identified and filtered out, information that does not correspond to reality may be more easily disseminated through education.

The second study focuses on the impact of using large language models (LLMs) versus traditional search engines for information gathering during learning, specifically examining how these tools affect cognitive load and learning outcomes (Stadler et al., 2024). In the study, 91 university students are divided into two groups. One group uses ChatGPT to research a topic and generate recommendations, while the other uses the search engine Google. As is well known, search engines direct users to relevant websites related to the topic they are investigating. This requires the user to extract, evaluate, and analyze the information they find, and ultimately construct a meaningful text within the context of their research. In contrast, LLMs take over this burden and generate answers directly, relieving the user of these cognitive tasks. In other words, when using a search engine, the individual is more actively engaged in the content, whereas with LLMs, their role becomes comparatively more passive. This distinction raises important questions: How do these two different approaches to information retrieval affect the quality of the content produced? And how do they differ in terms of the cognitive load placed on the user?

The findings of the study indicate that students who used LLMs experienced lower cognitive load across all three dimensions—extraneous, intrinsic, and germane—compared to those who used traditional search engines. The largest difference was observed in germane cognitive load, while the smallest was in extraneous cognitive load. This suggests that students using traditional search engines engaged in deeper processing and relied more heavily on germane cognitive resources when constructing a well-reasoned text. In short, because LLMs offer quick and easy access to information, they reduce cognitive load and lead students to invest less effort in processing information, resulting in more superficial engagement. As the authors emphasize, while LLMs lower mental effort, traditional search tasks—being more demanding—encourage students to think more deeply and actively process information into meaningful structures, thereby contributing more effectively to long-term learning.

On the other hand, the study's findings on the diversity of students' reasoned proposals show that those using traditional search engines were able to generate a significantly greater number of valid arguments. Since LLMs reduce cognitive load and minimize human contribution, the variety of justifications also declines. These findings suggest that while LLMs are effective at providing students with information, they impair students' ability to process that information critically and transform it into structured reasoning. In other words, easy access to information reduces both self-regulation and deep thinking. The study clearly shows that the negative impact of LLM use on the quality of reasoning stems from a lack of cognitive effort—specifically, low germane cognitive load. Therefore, systems like LLMs may substitute for students' self-regulatory skills, potentially bypassing essential processes such as planning, monitoring, and evaluation.

Finally, the last and most comprehensive study aims to reveal the cognitive cost of using large language model (LLM) tools—such as ChatGPT—in educational settings, particularly during essay writing (Kosmyna et al., 2025). Participants in the study were divided into three distinct groups for essay writing: the LLM group, which used a large language model; the search engine group, which relied on conventional internet searches; and the brain-only group, which used no external tools. Each group participated in three writing sessions. In a fourth and final session, the conditions were reversed: participants in the LLM group were asked to write an essay without using any tools, while those in the brain-only group were instructed to use an LLM for the first time. Throughout all sessions, participants' brain activity was recorded.

As expected, the group that used LLMs showed the least engagement of brain regions involved in effortful thinking, since they depended heavily on the AI tool. In

contrast, those using search engines were not only gathering information but also processing and integrating it into new knowledge, leading to higher levels of brain activity than the LLM group. The brain-only group, which relied solely on internal cognitive resources, was expected to exhibit the highest level of brain activation. Accordingly, the study seeks to answer several questions: How does brain activity differ when individuals write essays using an LLM, a search engine, or only their own cognitive effort? What qualitative differences emerge among the essays produced by these different groups? How is memory affected throughout the process? And does the sense of ownership over the produced text vary depending on the method used?

The findings of the comprehensive study reveal that while each group exhibited distinct neural connectivity patterns, participants within the same group showed consistent internal patterns. As expected, the extent of brain connectivity varied according to the intensity of brain engagement. In the LLM group—where cognitive effort was minimal—the weakest overall neural connections were observed. Conversely, in the brain-only group, which required the greatest mental effort, the strongest and most extensive neural networks were activated. Moreover, the essays produced by participants in the brain-only group demonstrated diverse approaches to each topic, while those written by participants in the LLM group were statistically more homogeneous and exhibited significantly less variation across topics. In other words, LLM tools not only result in lower brain activity but also increase uniformity in the output, thereby reducing the diversity and richness of the written content.

In the fourth session, although the LLM group—now writing without any tool—exhibited relatively higher neural connectivity compared to previous sessions, their prolonged prior exposure to LLMs resulted in less coordinated brain activity. Additionally, they showed a marked bias toward the vocabulary specific to LLM-generated language. On the other hand, the brain-only group, now using an LLM for the first time, demonstrated higher levels of memory recall during the writing process. Furthermore, participants in this group displayed stronger neural connectivity than those in the LLM group who had used the same tool during sessions 1, 2, and 3.

On the other hand, participants' sense of ownership toward their written essays also varied—as expected—based on the amount of effort they invested, which corresponded to the level of brain activity. According to the study's findings, the highest sense of ownership was observed among participants in the brain-only group, who wrote without using any external tools and therefore exerted the most mental effort. Participants in the LLM group not only reported the lowest levels of ownership but also showed a remarkably poor ability to recall or quote from the essays they had written just shortly before. In other words, participants in the LLM

group experienced the highest degree of alienation from the content they produced (Sidorkin, 2025). The reduction in cognitive load associated with LLM use appears to weaken memory retention. It is well established that learning tasks involving greater cognitive demands lead to higher levels of constructive cognitive load and, ultimately, better learning outcomes (Stadler et al., 2024). Therefore, this sense of alienation is directly linked to the level of brain engagement during the task. The involvement of generative AI in content creation processes undermines the gains that individuals typically acquire through direct, lived experience. As emphasized by the authors of the study discussed above, users of LLMs like ChatGPT tend to engage with content only at a superficial level. This diminishes the developmental cognitive load necessary for fostering active critical thinking.

The same characteristic applies for artistic production. In the case of artists, deeply complex and challenging experiences often lead them to a transformative stage in their creative journey. Their works bear witness to this process of becoming. While AI can replicate outputs that resemble those produced through such emotionally charged experiences, it severs the connection to the underlying emotions and blocks the experiential growth that the artist would gain through the act of creation. In this way, AI short-circuits the feedback loop between artistic struggle and personal development. As seen in the case of participants who relied solely on their own cognitive effort, the essential element in this process is the human being—particularly the artist—being on a path of becoming, continuously developing their creativity through lived experience. An artist's creations (or, in the context of the study above, the essays written by participants) serve as living witnesses to this ongoing process of personal growth. This is why artists tend to have a strong sense of ownership over their work: each creation corresponds to a challenging experience or a meaningful learning process. Through their works, we witness the story of an artist's becoming. However, this dimension is increasingly compromised with the use of artificial intelligence technologies. When AI is involved in the creation process, the human no longer undergoes a formative experience. Instead, they are reduced to the role of an operator issuing commands. In other words, the human-AI relationship remains confined to a technical level, and the person ultimately becomes alienated from the product that is generated.

In summary, three different studies mentioned in this Letter suggest that while generative AI applications may reduce cognitive load in the short term, they can weaken critical thinking and lead to cognitive decline in the long run. This negative impact is not limited to text production alone. A significant transformation is also occurring in how individuals search for and engage with information. In traditional search approaches, individuals are responsible for analyzing and synthesizing the

information they retrieve. However, with the rise of AI tools, these cognitive tasks are increasingly outsourced to the technology itself. As a result, individuals are deprived of the opportunity to develop essential skills through these processes. In other words, people are being steered away from active participation in cognitive tasks toward the passive consumption of ready-made content. This shift weakens the development of the core cognitive abilities that make us human while simultaneously increasing dependency on external tools like AI.

On the other hand, while generative AI applications such as ChatGPT that utilize LLMs offer clear advantages, they also come with significant cognitive costs. In human-machine interaction, AI tools tend to position the user in a relatively passive role, and because they reduce especially germane cognitive load, the learning that occurs risks remaining superficial. The findings from the studies discussed above regarding the diversity of students' arguments suggest that, in the long term, students' abilities for critical thinking and independent problem solving may deteriorate. Although these tools reduce both intrinsic and extraneous cognitive load, the additional time gained in the process did not lead to the generation of higher-quality arguments compared to students using traditional methods. In other words, while LLMs appear to relieve students from the demanding cognitive and experiential processes that deepen and solidify learning, if this trend continues, it may negatively affect the development of essential cognitive skills over time. Therefore, it is crucial to develop approaches that not only support the use of LLMs but also ensure that the reasoning produced is as diverse and high-quality as that generated through traditional methods.

Considering that self-regulation is a foundational component of lifelong learning policies—an essential requirement in the age of artificial intelligence—the risk increases that individuals may become more vulnerable and less resilient in the face of profound transformations in the labor market. In this new context, the human capacity to adapt may be weakened. Therefore, it is essential that AI applications, particularly within education systems, be used in ways that actively promote cognitive engagement and encourage interaction with complex content. In other words, generative AI tools should not be evaluated solely based on the convenience they provide, but also in terms of how they are used and what kinds of cognitive processes they stimulate.

The growing dependence on these tools, particularly among young users, poses critical long-term risks for the future of nations. The overreliance that comes with such dependency does not only raise concerns about the accuracy or bias of the content produced by AI; more importantly, it gradually undermines young people's ability to solve problems and think critically without relying on external tools. This

trend brings to the forefront a serious long-term threat: the prospect of future generations—who constitute a nation’s most vital asset, its human capital—growing up lacking the fundamental and essential skills needed for independent thought and action. For this reason, raising awareness of the cognitive and developmental risks associated with AI tool use is crucial, especially during the early stages of education.

Instead of using these applications as a valuable support to help students fulfill their responsibilities and improve the quality of their work, using them as substitutes for the student may, in the long term, lead to highly negative behavioral outcomes such as laziness, lack of responsibility, and weakened decision-making abilities. This is precisely where the critical importance of responsible AI use emerges. The ethical and responsible use of AI in education can help mitigate these negative effects. In this context, students and teachers should be informed about the long-term deterioration in human skills that may result from replacing human effort with AI. Raising awareness about how these applications can be used to complement rather than replace humans — and the boundaries of such use — is becoming increasingly vital in education.

In sum, this study examines the cognitive consequences of increasing reliance on generative artificial intelligence, particularly among younger populations. Drawing on three recent empirical studies, the findings suggest that while AI tools such as large language models can reduce cognitive load and improve efficiency in the short term, they also risk diminishing critical thinking, problem-solving, and self-regulation over time. The shift from active engagement with information to passive consumption of AI-generated content reduces opportunities for individuals to develop essential cognitive skills. In educational contexts, students using AI tools may produce arguments of similar surface quality to those using traditional methods, but without the deeper cognitive engagement required for long-term learning and intellectual growth. This trend raises serious concerns about the future resilience and adaptability of individuals in rapidly evolving labor markets, where independent thinking and learning are critical. These findings call for urgent attention to the ethical and responsible integration of AI in education and beyond. Policies must promote AI use as a complement to human effort rather than a substitute. Educators and policymakers should prioritize approaches that actively stimulate cognitive engagement and ensure that human reasoning is not outsourced, but enhanced. Without such measures, the long-term overdependence on AI risks eroding the very capacities that underpin individual autonomy, societal resilience, and the development of human capital.

References

- Ahmad, S. F., Han, H., Alam, M. M., Rehmat, M. K., Irshad, M., Arraño-Muñoz, M., & Ariza-Montes, A. (2023). Impact of artificial intelligence on human loss in decision making, laziness and safety in education. *Humanities and Social Sciences Communications*, 10:311. <https://doi.org/10.1057/s41599-023-01787-8>
- İlikhan, S., Özer, M., Tanberkan, H., & Bozkurt, V. (2024). How to mitigate the risks of deployment of artificial intelligence in medicine? *Turkish Journal of Medical Science*, 54(3), 483-492. <https://doi.org/10.55730/1300-0144.5814>
- İlikhan, S., Özer, M., Perc, M., Tanberkan, H., & Ayhan, Y. (2025). Complimentary use of artificial intelligence in healthcare. *Medical Journal of Western Black Sea*, 9(1), 7-17. <https://doi.org/10.29058/mjwbs.1620035>
- Kosmyna, N., Hauptmann, E., Yuan, Y. T., Situ, J., Liao, X. H., Beresnitsky, A. V., Braunstein, I., & Maes, P. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. <https://doi.org/10.48550/arXiv.2506.08872>
- Özer, M. (2024a). Potential benefits and risks of artificial intelligence in education. *Bartın University Journal of Faculty of Education*, 13(2), 232-244. <https://doi.org/10.14686/1416087>
- Özer, M. (2024b). Impact of ChatCPT on Scientific Writing. *The Journal of Humanity and Society*, 14(3), 210-217. <https://doi.org/10.12658/E0002>
- Özer, M. (2024c). Is artificial intelligence hallucinating?. *Turkish Journal of Psychiatry*, 35(4), 333-335. <https://doi.org/10.5080/u27587>
- Özer, M., & Perc, M. (2024). Human complementation must aid automation to mitigate unemployment effects due to AI technologies in the labor market. *Reflektif Journal of Social Sciences*, 5(2), 503-514. <https://doi.org/10.47613/reflektif.2024.176>
- Özer, M., Perc, M., & Suna, H. E. (2024a). Artificial intelligence bias and the amplification of inequalities in the labor market. *Journal of Economy, Culture and Society*, 69, 159-168. <https://doi.org/10.26650/JECS2023-1415085>
- Özer, M., Perc, M., & Suna, H. E. (2024b). Participatory management can help AI ethics adhere to the social consensus. *Istanbul University Journal of Sociology*, 44(1), 221-238. <https://doi.org/10.26650/SJ.2024.44.1.0001>
- Perc, M., Özer, M., & Hojnik, J. (2019). Social and juristic challenges of artificial intelligence. *Palgrave Communications*, 5(61). <https://doi.org/10.1057/s41599-019-0278-x>
- Sidorkin, A. M. (2025). Embracing liberatory alienation: AI will end us, but not in the way you may think. *AI & Society*, 40, 1417-1424. <https://doi.org/10.1007/s00146-024-02019-6>
- Stadler, M., Bannert, M., & Sailer, M. (2024). *Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry*. *Computers in Human Behavior*, 160:108386. <https://doi.org/10.1016/j.chb.2024.108386>
- Suna, H. E., & Özer, M. (2025). The human complimentary usage of AI and ML for fair and unbiased educational assessments. *Chinese/English Journal of Educational Measurement and Evaluation*, 6(1), 1-21. <https://doi.org/10.59863/YPKL4338>
- Tanberkan, H., Özer, M., & Gelbal, S. (2024). Impact of artificial intelligence on assessment and evaluation approaches in education. *International Journal of Educational Studies and Policy*, 5(2), 139-152. <https://doi.org/10.5281/zenodo.14016103>