

Is AI Making Us Stupid?

¹Trent N. Cash (tcash@uwaterloo.ca)

²Megan O. Kelly (mo.kelly@utoronto.ca)

³Brooke N. Macnamara (bmacnama@purdue.edu)

¹Evan F. Risko (efrisko@uwaterloo.ca)

¹Department of Psychology, University of Waterloo, Canada

²Department of Psychology, University of Toronto, Canada

³Department of Psychological Sciences, Purdue University, USA

This is an original manuscript of an article published in *Trends in Cognitive Sciences*, available at: <https://doi.org/10.1016/j.tics.2026.06.004>

Author Note

Correspondence concerning this article should be addressed to Trent N. Cash, Department of Psychology, University of Waterloo, 200 University Ave. W, PAS 4042, Waterloo, ON, N2L 3G1. E-Mail: tcash@uwaterloo.ca. Trent N. Cash was supported by the Lupina Foundation. Megan O. Kelly was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) and Evan F. Risko was supported by a NSERC Discovery Grant (#2025-02684). The authors report no competing interests.

Abstract

Many worry that AI use will erode human cognition. We discuss evidence suggesting that offloading cognition to AI can impede skill acquisition and lead to skill decay, but risks depend on how AI is used. We also consider whether our basic cognitive abilities might prove more resilient to this erosion.

Keywords: Artificial Intelligence; Cognitive Offloading; Skill Acquisition; Skill Decay; Human-AI Interaction

Cognitive Offloading

Our biological cognitive systems can be both taxing and fallible. As a result, humans have long sought support in pursuing our cognitive goals. For example, we have calculators to support arithmetic; GPS to support wayfinding; and the internet to support access to and storage of information. Often, using these supports reduces the cognitive labor that otherwise would be our cognitive system's obligation – a behavior generally referred to as *cognitive offloading* [1]. While the reduced cognitive demand and increased reliability from offloading are typically welcomed, this is often coupled with concern that offloading will compromise our cognitive functioning. This motif is captured well in Nicholas Carr's famous article "Is Google making us stupid?"¹ but has a long history spanning a wide variety of technologies and cognitive capacities. Recently, the proliferation of Large Language Models (LLMs), a class of artificial intelligence (AI) tools trained on human language and designed to assist with cognitive tasks (e.g., Claude, ChatGPT), has reignited conversations about the consequences of offloading. As such, many are asking: Will AI make us stupid?

The concern that AI will compromise our cognitive functioning appears largely driven by the intuition that offloading cognitive work to an AI will leave our cognitive systems less capable of performing that work without AI assistance. In addressing this concern, it is helpful to distinguish between two facets of our cognitive systems: skills and basic cognitive abilities. Skills are learned behaviors, supported by relevant knowledge (see Box 1 for discussion of AI's impact on acquiring knowledge), that are carried out with the aim of accomplishing a task effectively (e.g., doing arithmetic, flying airplanes). Skill acquisition has been studied for over a century, and successful theories have focused on how skills are acquired and improved with practice [2]. Basic cognitive abilities, on the other hand, can be thought of as foundational,

domain-general mental capacities (e.g., working memory, selective attention) that underlie the acquisition and execution of domain-specific skills.

Skill Acquisition and Decay

Given that skills are acquired through practice, offloading said practice to AI will almost certainly compromise skill acquisition [3]. For example, a recent study tasked high school students with learning a new mathematical concept either without AI access or with access to one of two different AI tools [4]. One of those tools allowed students to completely offload solving the practice problems. Students with access to this AI tool performed better on practice questions than students without AI support. However, on a subsequent test where no one had access to AI, students with prior access to this AI tool performed worse than students who never had the AI support. These results suggest that using this AI impeded their learning by preventing them from engaging in the practice needed to acquire the skill.

Like skill acquisition, practice also plays an important role in skill retention [3]. In this case, relying on an AI might cause an acquired skill to *decay* [3] because engaging in the cognitive work associated with that skill represents a necessary defense against forgetting [5]. For example, a recent study retrospectively evaluated the rate at which endoscopists detected the presence of polyps before and after an AI detection tool became available [6]. In the 3 months before the AI was introduced, physicians detected polyps in approximately 28.4% of patients. For the next 3 months, the AI tool was introduced to all physicians, with access to it randomly provided for some cases, but not others. In cases for which the AI was unavailable, detection rates fell to 22.4%. Thus, compared with their detection rates before the AI was introduced, the endoscopists' detection rates dropped when they could no longer access the AI. This suggests that the availability of the AI may have led to skill decay.

Basic Cognitive Abilities

While the concern that offloading to AI could impair the acquisition and retention of skills is well-grounded, it is less clear whether (or to what extent) offloading to AI will erode our basic cognitive abilities. One argument against the latter is that our basic cognitive abilities appear stubbornly resistant to substantial change. For example, research on cognitive training has generally found that while performance on cognitive tasks can be improved with training, this improvement is highly task-specific [7], rather than reflecting the kind of broad improvements in cognitive performance we might expect from changes to our basic cognitive abilities. Indeed, a recent review of research on cognitive training concluded that any limited training-induced cognitive change that does happen seems to be primarily through changes in how efficiently cognitive resources are used rather than the amount of cognitive resources available to the individual [8]. Thus, while the learned skills that we offload to AI may have a dubious future, the basic cognitive abilities on which these skills are based may be less likely to share this fate (or not to the same extent at least; see also [9]). That said, there are certain factors (e.g., schooling) that appear to be associated with modest changes in basic cognitive abilities [10], and it remains an open question whether long-term offloading to AI, or offloading at different developmental stages (e.g., early childhood), may have a similar impact.

How We Use AI

Whether our skills remain intact when offloading to AI will depend, in part, on the form the offloading takes (e.g., see [11] for relevant work on levels of automation). For example, imagine a student with a math problem to solve. That student might have the AI submit an answer on their behalf, offloading the cognitive work entirely. Alternatively, the student might have the AI answer the problem but then review the solution provided by the AI. The student

might also have the AI provide an explanation in addition to the solution or merely ask the AI to provide a suggestion about how to approach the problem. In each of these cases (and one could imagine many more), the AI's impact on cognition is likely to be different.

Indeed, recent and emerging work investigating the impact of AI supports this claim. In the mathematical learning study with high school students described earlier [4], the authors also included a third condition in which students used a *custom* AI tutor designed to probe their knowledge and fill gaps in their learning (akin to a thoughtful tutor), rather than giving them the answers. Unlike the earlier-mentioned students who used the AI that allowed them to completely offload solving the practice problems, these students performed no worse than those without access to AI when the AI was taken away, suggesting that they still developed the targeted mathematical skill. There is also preliminary evidence that using AI as a collaborator that provides feedback [12] or using AI's outputs as examples to be emulated [13] can benefit future performance, even in cases when the AI is no longer available. Individuals may also get better at leveraging these kinds of benefits over time, as they develop skills, such as prompt engineering, that allow them to work with AI more effectively [12]. Taken together, these results suggest that how we offload to, and engage with, AI will likely determine its impact on cognition.

Concluding Remarks

As AI-based technologies advance and continue to provide cognitive support in novel ways, concerns about the future of cognition warrant careful reflection. Will AI make us stupid? It is far too early to say with certainty what the long-term effects of offloading cognitive work to AI will be, and there are many open questions about the impacts that offloading may have on our cognitive functioning (see Table 1). However, basic and applied research in cognitive science and allied fields (e.g., human factors research, learning science), and emerging work on AI's impacts

on cognition, provide credible expectations for what lies ahead. First, we can expect that offloading relevant cognitive work entirely to AI will hinder the acquisition, or foster the decay, of learned skills. But these outcomes are not inevitable. If we remain in the proverbial cognitive loop when using AI tools, we can mitigate the costs of offloading and may even be able to structure our interactions with AI to boost the acquisition and/or retention of skills.

Furthermore, the costs of offloading may be limited to the skills that are specifically offloaded to AI, rather than broadly affecting our basic cognitive abilities. While we should not be too cavalier about losing such skills, this may provide something of a bulwark against more dystopian cognitive futures in which AI undermines our most fundamental cognitive capacities. In sum, there is clearly risk that AI can make us “stupid” by compromising our skills (and knowledge) if we completely offload them to AI. On the other hand, AI may be less likely to diminish the foundational cognitive capacities that make us smart, not *stupid*, in the first place.

Resources

¹<https://www.theatlantic.com/magazine/archive/2008/07/is-google-making-us-stupid/306868/>

References

- [1] Risko, E. F., & Gilbert, S. J. (2016). Cognitive offloading. *Trends in Cognitive Sciences*, 20(9), 676-688. <https://doi.org/10.1016/j.tics.2016.07.002>
- [2] Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Brooks/Cole.
- [3] Macnamara, B. N., Berber, I., Çavuşoğlu, M. C., Krupinski, E. A., Nallapareddy, N., Nelson, N. E., Smith, P. J., Wilson-Delfosse, A. L., & Ray, S. (2024). Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers' awareness? *Cognitive Research: Principles and Implications*, 9, Article 46. <https://doi.org/10.1186/s41235-024-00572-8>
- [4] Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabacki, O., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122. <https://doi.org/10.1073/pnas.2422633122>
- [5] Bjork, R. A., & Bjork, E. L. (1992). A new theory of disuse and an old theory of stimulus fluctuation. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *From Learning Processes to Cognitive Processes: Essays in honor of William K. Estes* (Vol. 2, pp. 35-67). Hillsdale, NJ: Erlbaum.
- [6] Budzyń, K., Romańczyk, M., Kitala, D., Kołodziej, P., Bugajski, M., Adami, H. O., Blom, J., Buszkiewicz, M., Halvorsen, N., Hassan, C., Romańczyk, T., Holme, Ø., Jarus, K., Fielding, S., Kunar, M., Pellisé, M., Pilonis, N., Kamiński, M. F., Kalager, M., Bretthauer, M., & Mori, Y. (2025). Endoscopist deskilling risk after exposure to artificial intelligence

- in colonoscopy: A multicentre, observational study. *The Lancet Gastroenterology & Hepatology*, 10(10), 896–903. [https://doi.org/10.1016/S2468-1253\(25\)00133-5](https://doi.org/10.1016/S2468-1253(25)00133-5)
- [7] Sala, G., & Gobet, F. (2019). Cognitive training does not enhance general cognition. *Trends in Cognitive Sciences*, 23(1), 9–20. <https://doi.org/10.1016/j.tics.2018.10.004>
- [8] von Bastian, C. C., Belleville, S., Udale, R. C., Reinhartz, A., Essounni, M., & Strobach, T. (2022). Mechanisms underlying training-induced cognitive change. *Nature Reviews Psychology*, 1, 30–41. <https://doi.org/10.1038/s44159-021-00001-3>
- [9] Pearson, H. (2025). Are the Internet and AI affecting our memory? What the science says. *Nature*, 638(8049), 26–28. <https://doi.org/10.1038/d41586-025-00292-z>
- [10] Ritchie, S. J., & Tucker-Drob, E. M. (2018). How much does education improve intelligence? A meta-analysis. *Psychological Science*, 29(8), 1358–1369. <https://doi.org/https://doi.org/10.1177/0956797618774253>
- [11] Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 30(3), 286–297. <https://doi.org/10.1109/3468.844354>
- [12] Oppenheimer, D. M., Cash, T. N., & Connell Pensky, A. E. (2025). You’ve got AI friend in me: LLMs as collaborative learning partners. *International Journal of Artificial Intelligence in Education*, 35(6), 3896–3921. <https://doi.org/10.1007/s40593-025-00521-6>

- [13] Lira, B., Rogers, T., Goldstein, D. G., Ungar, L., & Duckworth, A. L. (2025). *Coach not crutch: Evidence that AI can improve writing skill despite reducing effort* [Preprint]. ArXiv. <https://doi.org/10.48550/arXiv.2502.02880>
- [14] Kelly, M. O., & Risko, E. F. (2022). Study effort and the memory cost of external store availability. *Cognition*, 228, 105228. <https://doi.org/10.1016/j.cognition.2022.105228>
- [15] Melumad, S., & Yun, J. H. (2025). Experimental evidence of the effects of large language models versus web search on depth of learning. *PNAS Nexus*, 4(10), pgaf316. <https://doi.org/10.1093/pnasnexus/pgaf316>

Box 1: Acquiring Knowledge (or not) with AI

While we focus on skills in the main text, offloading to AI may also impact our acquisition of knowledge. We define acquired knowledge as information stored in and retrievable from memory that was obtained through instruction, study, or another form of experience. There is considerable evidence demonstrating that when people can offload to an external store (e.g., computer) during learning, doing so reduces the likelihood they will later be able to retrieve the offloaded information from memory [14]. This appears to be at least partly due to individuals investing less effort during learning. A recent study demonstrated a similar pattern with AI. Participants were tasked with learning about a topic using either Google search or an AI [15]. Participants then had to give advice to a friend based on the information they learned, without technological assistance. Advice given by participants in the AI condition was briefer, less unique to the writer, and rated as less helpful. In addition, participants who used the AI spent less time learning, learned less information, and felt less ownership over the knowledge they acquired. The authors argued that Google searches required learners to open links and summarize information across sources, whereas AI tools summarized the information on the learner's behalf. Thus, these AI summaries allowed participants to offload the cognitive work related to information search and summarization, leading to a shallower depth of learning. Hence, offloading to an AI during learning can compromise knowledge acquisition. That said, AI also has the potential to benefit knowledge acquisition in a variety of ways (e.g., finding information you may not have found otherwise).

Table 1: Open Questions about Offloading to AI

Does the impact of offloading to AI vary as a function of one's developmental stage (e.g., children, younger adults, older adults)?
Are there key developmental phases where offloading to AI may be particularly harmful?
How does prolonged offloading to AI impact cognition (e.g., over years or decades)?
How does the level of initial skill acquisition (before the introduction of AI) impact decay induced by offloading?
How might we use refresher training or other strategies to maintain skills in environments where offloading to AI is typical?
Can individuals learn to strategically offload to AI, minimizing its cognitive costs while maximizing its benefits?
How does offloading to AI impact our metacognitive skills?
How can AI be designed to discourage more harmful forms of offloading?