

Technology

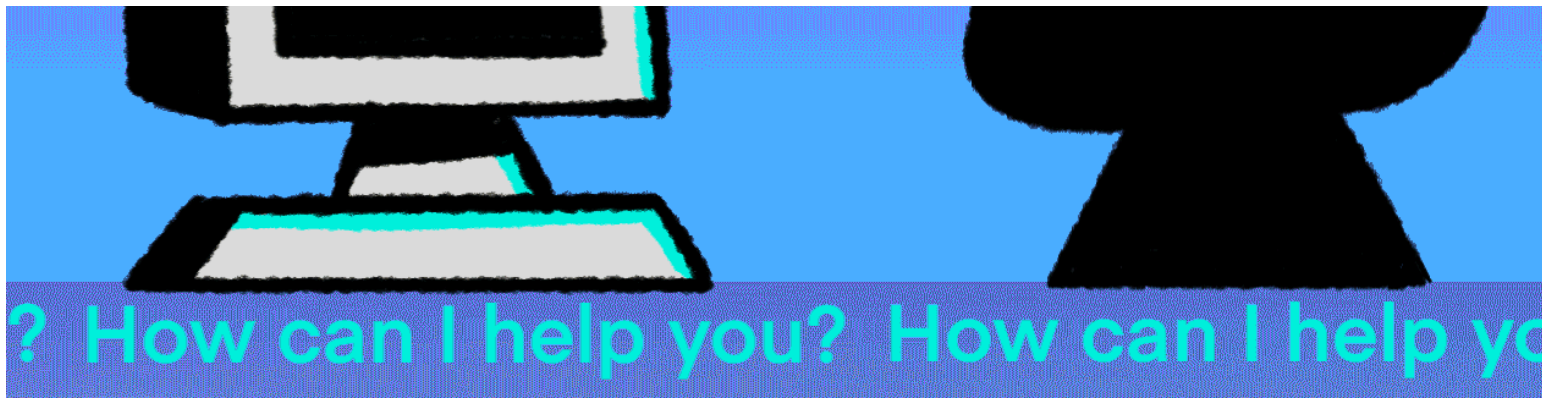
AI may blunt our thinking skills – here's what you can do about it

There is growing evidence that our reliance on generative AI tools is reducing our ability to think clearly and critically, but it doesn't have to be that way

By [Miriam Frankel](#)

📅 10 November 2025





👤 Allan Sanders

Socrates wasn't the greatest fan of the written word. Famous for leaving no texts to posterity, the great philosopher is said to have believed that a reliance on writing [destroys the memory and weakens the mind](#) 🔗

<https://blogs.ubc.ca/etec540sept13/2013/09/29/socrates-writing-vs-memory/>.

Some 2400 years later, Socrates's fears seem misplaced – particularly in light of evidence that [writing things down improves memory](#) 🔗 [/article/2414241-writing-things-down-may-help-you-remember-information-more-than-typing/](#). But his broad mistrust of cognitive technologies lives on. A growing number of psychologists, neuroscientists and philosophers worry that ChatGPT and similar generative AI tools will chip away at our powers of information recall and blunt our capacity for clear reasoning.

What's more, while Socrates relied on clever rhetoric to make his argument, these researchers are grounding theirs in empirical evidence. Their studies have uncovered evidence that even trained professionals disengage their critical thinking skills when using generative AI and revealed that an over-reliance on these AI tools during the learning process reduces brain connectivity and renders information less memorable. Little wonder, then, that when I asked Google's Gemini chatbot whether AI tools are turning our brains to jelly and our memories to sieves, it admitted they might be. At least, I think it did: I can't quite remember now.

But all is not lost. Many researchers suspect we can flip the narrative, turning generative AI into a tool that improves our cognitive performance and augments our intelligence. "AI is not necessarily making us stupid, but we may be interacting with it stupidly," says [Lauren Richmond](#) 🔗 https://www.stonybrook.edu/commcms/psychology/faculty/_faculty_profiles/lrichmond at Stony Brook University, New York. So, where are we going wrong with generative AI tools? And how can we change our habits to make better use of the technology?

The generative AI age

In recent years, generative AI has become deeply embedded in our lives. [Therapists use it to look for patterns](#) 🔗 <https://www.apaservices.org/practice/news/artificial-intelligence-psychologists-work-in-their-notes>. Students [rely on it for essay writing](#) 🔗 [/article/2354663-should-schools-ban-chatgpt-or-embrace-the-technology-instead/](#). It has even been welcomed by media organisations, which may be why financial news website Business Insider reportedly [now permits its journalists to use AI](#) 🔗 <https://www.theverge.com/news/779739/business-insider-ai-writing-stories> when drafting stories.

In one sense, all of these AI users are following a millennia-old tradition of "cognitive offloading" – using a tool or physical aid to reduce mental burden. Many of us use this strategy in our daily lives. Every time we write a shopping list instead of memorising items to buy, we are employing cognitive offloading.

Used in this way, cognitive offloading can help us improve our accuracy and efficiency, while simultaneously freeing up brain space to handle more complex cognitive tasks such as problem-solving, says Richmond. But in a review of the behaviour that Richmond published earlier this year with her Stony Brook colleague [Ryan Taylor](#) 🔗 <https://www.linkedin.com/in/ryan-taylor-95b37b292>, they found [it has negative effects on our cognition too](#) 🔗 <https://www.nature.com/articles/s44159-025-00432-2>.

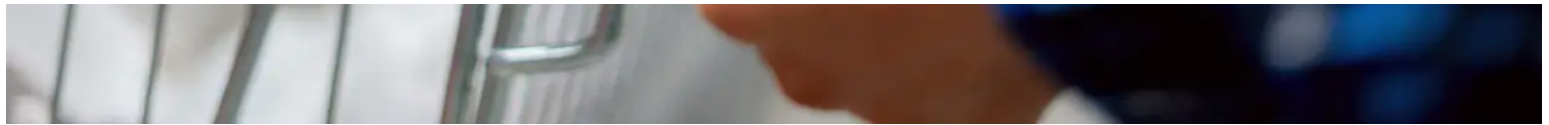
“When you’ve offloaded something, you almost kind of mentally delete it,” says Richmond. “Imagine you make that grocery list then you don’t take it with you. You’re actually worse off than if you just planned on remembering the items that you needed to the store.”

Research backs this up. To take one example, a study published in 2018 revealed that when we take photos of objects we see during a visit to a museum, we are [worse at remembering what was on display afterwards](#) [🔗](#) <https://www.sciencedirect.com/science/article/abs/pii/S002210311730505X>: we have subconsciously given our phones the task of memorising the objects on show.

This can create a spiral whereby the more we offload, the less we use our brains, which in turn makes us offload even more. “Offloading begets offloading – it can happen,” says [Andy Clark](#) [🔗](#) <https://profiles.sussex.ac.uk/p493-andy-clark>, a philosopher at the University of Sussex, UK. In 1998, Clark and his colleague [David Chalmers](#) [🔗](#) <https://profiles.sussex.ac.uk/mg25333710-900-david-chalmers-interview-virtual-reality-as-real-as-real-reality/> – now at New York University – proposed the [extended mind thesis](#) [🔗](#) <https://www.jstor.org/stable/33> which argues that our minds extend into the physical world through objects such as shopping lists and photo albums. Clark doesn’t see that as inherently good or bad – although he is concerned that as we extend into cyberspace with generative AI and other online services, we are making ourselves vulnerable if those services ever become unavailable because of power cuts or cyberattacks.

Cognitive offloading could also make our memory more vulnerable to manipulation. In a 2019 study, researchers at the University of Waterloo, Canada, presented volunteers with a list of words to memorise and allowed them to type out the words to help remember them. The researchers found that when they [surreptitiously added a rogue word to the typed list](#) [🔗](#) <https://pubmed.ncbi.nlm.nih.gov/31330472/>, the volunteers were highly confident that the rogue word had actually been on the list along.





🚩 We cognitively offload whenever we write a shopping list

Mikhail Rudenko/Alamy

As we have seen, concerns about the harms of cognitive offloading go back at least as far as Socrates. But generative AI has supercharged them. In a study posted online this year, [Shiri Melumad](http://www.shirimelumad.com/) and [Jin Ho Yun](https://drjyun.github.io/) at the University of Pennsylvania asked 1100 volunteers to write a short essay offering advice on planting a vegetable garden after researching the topic either using a standard web search or ChatGPT. The resulting essays tended to be shorter and contained fewer references to facts if they were written by volunteers who used ChatGPT, which the researchers interpreted as evidence that the AI tool had made the learning process more passive – and the resulting understanding more superficial. Melumad and Yun argued that this is because the AIs synthesise information in other words, we cognitively offload our opportunity to explore and make discoveries about a subject for ourselves.

Sliding capacities

The latest neuroscience is adding weight to these fears. In experiments detailed in [a paper](https://arxiv.org/abs/2506.08872) peer review which was released this summer, [Nataliya Kos'myna](https://www.media.mit.edu/people/nkosmyna/overview/) at Massachusetts Institute of Technology and her colleagues used EEG head caps to measure the brain activity of 54 volunteers as they wrote essays on subjects such as “Does true loyalty require unconditional support?” and “Is having too many choices a problem?” of the participants wrote their essays using just their own knowledge and experience, those in a second group were allowed to use a Google search engine to explore the essay subject, and a third group could use ChatGPT.

The team discovered that the group using ChatGPT had lower brain connectivity during the task while the group relying simply on their own knowledge had the highest. The browser group, meanwhile, was somewhere in between.

“There is definitely a danger of getting into the comfort of this tool that can do almost everything. And that can have a cognitive cost,” says Kos'myna.

Critics may argue that a reduction in brain activity needn't indicate a lack of cognitive involvement in an activity, which Kos'myna accepts. “But it is also important to look at behavioural measures,” she says. For example, when quizzing the volunteers later, her colleagues discovered that the ChatGPT users found it harder to quote their essays, suggesting they hadn't been as invested in the writing process.

There is also emerging – if tentative – evidence of a link between heavy generative AI use and poorer critical thinking. For instance, [Michael Gerlich](https://www.sbs.edu/team/dr-michael-gerlich/) at the SBS Swiss Business School published a study earlier this year

assessing the AI habits and critical thinking skills of 666 people from diverse backgrounds.

Gerlich used structured questionnaires and in-depth interviews to quantify the participants' critical thinking skills, which revealed that those aged between 17 and 25 had critical thinking scores that were roughly 45 per cent lower than participants who were over 40 years old.



⚠️ We remember less of what we see when we use our cameras
Grzegorz Czapski/Alamy

“These [younger] people also reported that they depend more and more on AI,” says Gerlich: they were between 40 and 45 per cent more likely to say they relied on AI tools than older participants. In combination, Gerlich thinks the two findings hint that [over-reliance on AI reduces critical thinking skills](https://www.mdpi.com/2075-4698/15/1/6#:~:text=Cognitive%20offloading%20refers%20to%20the,involvement,%20potentially%20affecting%20critical%20) <https://www.mdpi.com/2075-4698/15/1/6#:~:text=Cognitive%20offloading%20refers%20to%20the,involvement,%20potentially%20affecting%20critical%20>

Others stress that it is too early to draw any firm conclusions, particularly since Gerlich’s study showed correlation rather than causation – and given that some research suggests [critical thinking skills are inherently underdeveloped in adolescents](https://eric.ed.gov/?id=EJ1301306) <https://eric.ed.gov/?id=EJ1301306>. “We don’t have the evidence yet,” says [Aaron French](https://www.linkedin.com/in/dr-french/) <https://www.linkedin.com/in/dr-french/> at Kennesaw State University in Georgia.

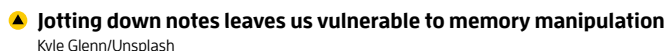
But other research suggests the link between generative AI tools and critical thinking may be real. In a study published earlier this year by a team at Microsoft and Carnegie Mellon University in Pennsylvania, 319 “knowledge workers” (scientists, software developers, managers and consultants) were asked about their experiences with generative AI. The researchers found that people who expressed higher confidence in the technology [freely admitted to engaging in less critical thinking while using it](https://www.microsoft.com/en-us/research/wp-content/uploads/2025/01/lee_2025_ai_critical_thinking_survey.pdf) https://www.microsoft.com/en-us/research/wp-content/uploads/2025/01/lee_2025_ai_critical_thinking_survey.pdf. This fits with Gerlich’s suspicion that an over-reliance on AI tools instils a degree of “cognitive laziness” in people.


Perhaps most worrying of all is that generative AI tools may even influence the behaviour of people who don’t use the tools head-on. In a study published earlier this year, [Zachary Wojtowicz](https://www.zachary-wojtowicz.com/) <https://www.zachary-wojtowicz.com/> and [Simon DeDeo](https://www.cmu.edu/dietrich/sds/people/faculty/simon-dedeo.html) <https://www.cmu.edu/dietrich/sds/people/faculty/simon-dedeo.html> – who were both at Carnegie Mellon University at the time – argued that we have learned to value the effort that goes into certain behaviours, like craft and thoughtful and sincere apology in order to repair social relationships. If we can’t escape the suspicion that someone has offloaded cognitively tricky tasks onto an AI – having the technology draft an apology on their behalf, say – [we may be less inclined to be genuine](https://ojs.aaai.org/index.php/AAAI/article/view/32151) <https://ojs.aaai.org/index.php/AAAI/article/view/32151>.

Using tools intelligently

One way to avoid all of these problems is to reset our relationship with generative AI tools, using them in a way that enhances rather than undermines cognitive engagement. That isn’t as easy as it sounds. In a new study, Gerlich found that even volunteers who took care to monitor themselves on their critical thinking skills have a tendency to [slip into lazy cognitive habits when using generative AI tools](https://www.mdpi.com/2306-5729/10/11/172) <https://www.mdpi.com/2306-5729/10/11/172>. “As soon as they were using generative AI without guidance, most of them directly offloaded,” says Gerlich.


When there is guidance, however, it is a different story. Supplemental work by Kos’myrna and her colleagues provides a good example. They asked the volunteers who had written an essay using only their own knowledge to work on a second version of the same essay at the same time using ChatGPT to help them. The EEG data showed that these volunteers maintained high brain connectivity even as they used the AI tool.



French sees the benefit in this approach too. In a paper he published last year with his colleague, the late J.P. Shim, he argued the right way to think about generative AI is as a [tool to enhance your existing understanding of a given subject](https://link.springer.com/article/10.1007/s10796-024-10562-2)  <https://link.springer.com/article/10.1007/s10796-024-10562-2>. The wrong way, meanwhile, is to view the tool as a convenient that replaces the need for you to develop or maintain any understanding.

So what are the secrets to using AI the right way? Clark suggests we should begin by being a bit less trusting: “Treat it like a co that sometimes has great ideas, but sometimes is entirely off the rails,” he says. He also believes that the more thinking you do using a generative AI tool, the better what he dubs your “hybrid cognition” will be.

That being said, Clark says there are times when it is “safe” to be a bit cognitively lazy. If you need to bring together a lot of pu available information, you can probably trust an AI to do that, although you should still double-check its results.

Gerlich agrees there are good ways to use AI. He says it is important to be aware of the “anchoring effect” – a cognitive bias th us rely heavily on the first piece of information we get when [making decisions](#)  [/article/mg21228381-800-decision-time-how-forces-shape-your-choices/](#). “The information you first receive has a huge impact on your thoughts,” he says. This means that you think you are using AI in the right way – critically evaluating the answers it produces for you – you are still likely to be gui what the AI told you in the first place, which can serve as an obstacle to truly original thinking.

But there are strategies you can use to avoid this problem too, says Gerlich. If you are writing an essay about the French Revolu negative impacts on society, don’t ask the AI for examples of those negative consequences. “Ask it to tell you facts about the Fr Revolution and other revolutions. Then look for the negatives and make your own interpretation,” he says. A final stage might i sharing your interpretation with the AI and asking it to identify any gaps in your understanding, or to suggest what a counter-ε might look like.

The AI expert who says artificial general intelligence is nonsense


Artificial intelligence has more in common with ants than humans, says Neil Lawrence. Only by taking a more nuanced view of intelligence can we see how machines will truly society

 /article/mg26335091-000-the-ai-expert-who-says-artificial-general-intelligence-is-nonsense/

This may be easier or harder depending on who you are. To use AI most fruitfully, you should know your strengths and weaknesses. For example, if you are experiencing cognitive decline, then offloading may offer benefits, says Richmond. Personality could also play a role. If you enjoy thinking, it is a good idea to use AI to challenge your understanding of a subject instead of asking it to spoon-feed facts.

Some of this advice may seem like common sense. But Clark says it is important that as many people as possible are aware of it for a simple reason: if more of us use generative AI in a considered way, we may actually help to keep those tools sharp.

If we expect generative AI to provide us with all the answers, he says, then we will end up producing less original content ourselves. Ultimately, this means that the large language models (LLMs) that power these tools – which are trained using human-generated data – will start to decline in capacity. “You begin to get the danger of what some people call model collapse,” he says: the LLMs are fed into feedback loops where they are trained on their own content, and [their ability to provide creative, high-quality answers deteriorates](#).

 <https://www.ibm.com/think/topics/model-collapse>. “We’ve got a real vested interest in making sure that we continue to write useful and interesting things,” says Clark.