

The Skill You're Renting Back

Cognitive offloading and the quiet atrophy of the human mind.

Kymata Labs Research · An independent research institution studying how AI systems actually behave in~12 min read production.

Tags · Cognitive Science · AI & Work · Skill Atrophy · Human Capital

Every task we hand to AI is a skill we stop practising, and skills we stop practising, we lose. Drawing on a widely-cited MIT Media Lab EEG preprint, a Lancet study showing experienced endoscopists deskilled within three months of AI-assisted colonoscopy, decades of aviation-automation research, and the navigation and memory literature, this paper argues that for most knowledge workers AI is not augmenting human capability but quietly hollowing it out. The danger is not that the tools get too good. It is that we get worse, and do not notice until the tool is taken away.

The argument, in moves

1. **Offloading a task offloads the practice that keeps the skill.** A capability you stop exercising decays. This is not a metaphor; it is how skill memory works.
2. **AI offloads the general act of thinking in language.** A calculator takes arithmetic, a GPS takes navigation. A language model takes reasoning, drafting, weighing, deciding, the broadest cognitive load any tool has ever absorbed.
3. **The deskilling is already measured.** Clinicians lost measurable performance within months; knowledge workers report reduced critical-effort; students who leaned on the model early showed weaker recall and brain connectivity.
4. **The order of use is the variable that matters.** Effort first, then assistance, preserves the skill. Assistance first replaces it.
5. **The cost lands at the worst moment.** Automation concentrates the need for human judgement into the rare cases the tool gets wrong, exactly when your atrophied skill is required.

The one-liner: the danger isn't that AI gets too good. It's that we get worse, and don't notice until the tool isn't there.

The strongest evidence: a clinic, not a lab

The most robust result in this set is not from a psychology lab; it is from a hospital. A study in *The Lancet Gastroenterology & Hepatology* measured experienced endoscopists before and after they began using AI-assisted colonoscopy, and found their unaided detection performance degraded within roughly three months of exposure³. This is a real-world, multi-centre, before-and-after measurement of clinical skill, on professionals who were already expert. The AI did not make them better when it was switched off. It made them worse.

The brain on an assistant

A widely-cited MIT Media Lab preprint, “Your Brain on ChatGPT,” used EEG to study people writing essays with and without an AI assistant, and reported the accumulation of what the authors call *cognitive debt*: weaker neural connectivity and poorer recall among those who leaned on the model ¹. The sample is small and the work is not yet peer-reviewed, and we flag that. But its most useful finding is about sequence: participants who wrote unaided first and then brought in the AI retained stronger recall and connectivity than those who started with the model. The order, effort then assistance, is the whole game.

Microsoft Research and Carnegie Mellon, surveying knowledge workers at CHI 2025, found self-reported reductions in cognitive effort and shifts in confidence when generative AI was in the loop ². Gerlich (2025) reports a correlation between heavy AI-tool use and lower critical-thinking scores, mediated by cognitive offloading ⁶. These are correlational and self-reported, weaker than the colonoscopy result, but they point the same direction.

The pattern repeats across domains

What makes the thesis hard to dismiss is not any single study; it is that the same shape recurs across fields that share almost nothing else.

Domain	Finding	Source
Medicine	Endoscopists deskilled within ~3 months of AI assistance	Lancet G&H, 2025
Knowledge work	Self-reported drop in critical effort with genAI	Microsoft / CMU, CHI 2025
Aviation	Manual-flying skill decays in the automated cockpit	Casner et al., 2014; FAA SAFO
Navigation	Habitual GPS use degrades spatial memory	Dahmani & Bohbot, 2020
Memory	“Google effect”: we recall where, not what	Sparrow et al., Science 2011

Aviation learned this lesson first and at the highest stakes: the FAA has issued safety alerts urging pilots to hand-fly precisely because automation erodes the manual skill needed when it disengages ⁵. The navigation literature shows habitual GPS use measurably degrades the spatial memory that London taxi drivers famously over-developed ⁴. And the “Google effect” showed, back in 2011, that when we know a fact is retrievable, we remember where to find it rather than the fact itself ⁷. None of these tools is as broad as a language model.

Cognitive offloading, and the trade it makes

Offloading is not inherently bad; it is how civilisation scales. We offload memory to writing, arithmetic to calculators, direction to maps. The trade is favourable when the tool is narrow and the freed capacity goes to higher work. It turns unfavourable when the offloaded capability is the general one, thinking a problem through, and when the freed capacity goes nowhere. The narrower the tool, the safer the trade. A general-purpose reasoning engine is the least narrow tool ever built.

This is not the printing press

The obvious objection is that every generation panics: Socrates feared writing would ruin memory. He was wrong about civilisation and right about memory, and that is the point. What has changed is scope. Past tools offloaded a bounded faculty. A language model offloads the open-ended act of reasoning in language across nearly every knowledge domain at once. The historical pattern should make us humble about predicting collapse, and precise about what is actually being traded away.

The moment it matters

If the AI is always there, why does the lost skill matter? Because it matters in exactly the moment the tool is not there, is not right, or is not trusted: the failed instrument, the edge case, the confidently wrong answer. Automation does not remove the need for human judgement; it concentrates that need into rare, high-stakes moments, then quietly erodes your readiness for them. The skill you stopped practising is the one you will need when the system fails.

What to do: keep one hard thing unaided

The evidence speaks to *how* you use AI, not *whether*. Use it to extend a skill you already practise, not to replace the practice. Preserve the order: do the hard cognitive work first, then bring the model in to check, extend, and accelerate. And keep one hard thing unaided, a kind of writing, a clinical read, mental arithmetic, navigating your own city, practised often enough that it stays yours. Treat the AI as a power tool you reach for after the manual work, not as a replacement for the hand.

Frequently asked questions

Isn't this just the same panic people had about calculators and the printing press?

Partly, and that history should make us humble rather than complacent. Socrates worried writing would ruin memory; he turned out wrong about civilisation and right about memory. What has changed now is scope. A calculator offloads arithmetic and a GPS offloads navigation, but a large language model offloads the general-purpose act of thinking through a problem in language: reasoning, drafting, weighing, deciding. The narrower the tool, the safer the trade, and this one is anything but narrow.

Does this mean I shouldn't use AI?

No. The evidence speaks to how you use it, not whether you should. The MIT Media Lab preprint found that people who wrote unaided first and then brought in the AI kept stronger recall and brain connectivity than those who leaned on the model from the start. The order is what mattered: effort, then assistance. Use AI to extend a skill you already practise, rather than to replace the practice itself.

How strong is the evidence, really?

It varies, and we have flagged that inline. The colonoscopy finding in The Lancet is the strongest of the set: a real-world, multi-centre, before-and-after measurement of clinical performance. The MIT Media Lab paper is a widely-cited preprint with a small sample, not yet peer-reviewed. The survey and offloading studies are correlational. No single study proves the thesis on its own. What makes it hard to dismiss is the way the same pattern recurs across domains as different as medicine, aviation, navigation, and knowledge work.

If the AI is always available, does losing the underlying skill even matter?

It matters in exactly the moment the tool isn't there, isn't right, or isn't trusted: the failed instrument, the edge case, the answer that is confidently wrong. Automation doesn't remove the need for human judgement. It concentrates that need into rare, high-stakes moments, and then quietly erodes your readiness for them. The skill you stopped practising is the one you will need when the system fails.

What's the single most useful thing an individual can do?

Keep one hard thing unaided. Pick a skill that matters to you, whether it is writing, a clinical read, mental arithmetic, or navigating your own city, and practise it without assistance often enough that it stays yours. Treat the AI as a power tool you reach for after the manual work, not as a replacement for the hand.

References

- 1 Kosmyna, N. et al. (2025). "Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing." MIT Media Lab, arXiv:2506.08872 (preprint, not peer-reviewed). <https://arxiv.org/abs/2506.08872>
- 2 Lee, H.-P. et al. (2025). "The Impact of Generative AI on Critical Thinking." Microsoft Research & Carnegie Mellon University, CHI 2025. <https://www.microsoft.com/en-us/research/publication/the-impact-of-generative-ai-on-critical-thinking-self-reported-reductions-in-cognitive-effort-and-confidence-effects-from-a-survey-of-knowledge-workers>
- 3 Budzyń, K. et al. (2025). "Endoscopist deskilling after exposure to artificial intelligence in colonoscopy." The Lancet Gastroenterology & Hepatology. <https://pubmed.ncbi.nlm.nih.gov/40816301>
- 4 Dahmani, L. & Bohbot, V. D. (2020). "Habitual use of GPS negatively impacts spatial memory during self-guided navigation." Scientific Reports, 10, 6310. <https://www.nature.com/articles/s41598-020-62877-0>
- 5 Casner, S. M. et al. (2014). "The Retention of Manual Flying Skills in the Automated Cockpit." Human Factors, 56(8). (See also FAA SAFO 13002 and 17007.) <https://pubmed.ncbi.nlm.nih.gov/25509828>
- 6 Gerlich, M. (2025). "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking." Societies, 15(1), 6. <https://www.mdpi.com/2075-4698/15/1/6>
- 7 Sparrow, B., Liu, J. & Wegner, D. M. (2011). "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips." Science, 333(6043). <https://www.science.org/doi/10.1126/science.1207745>