Does Using AI Impact Critical Thinking and Creativity?

A Review of Emerging Evidence and Implications for Risk-Critical Industries

Author:
Dr Anya Keatley
Head of AI, Secret Compass

Date: December 2025

Executive Summary

As generative AI becomes embedded in everyday professional workflows, questions have emerged about its impact on human cognition. Recent studies from MIT, Microsoft, the British Journal of Educational Technology and Stanford/BetterUp Labs explore whether AI use reduces critical thinking, narrows creativity or encourages over-reliance that weakens long-term skill development.

Across these studies, a consistent pattern emerges. AI can increase speed and efficiency, but often shifts cognitive effort away from deep reasoning and towards verification and passive review. Users frequently apply less critical evaluation, retain less of what they write or read and generate a narrower range of ideas. Some researchers describe this as cognitive debt or metacognitive laziness.

This paper reviews these findings and discusses the related concept of workslop: AI-generated outputs that appear polished but lack depth or meaningful substance. Workslop increases the risk of shallow decision-making, particularly in environments where nuance and judgement matter.

The paper then explores hybrid intelligence as a framework for responsible AI use. Rather than replacing human thinking, hybrid systems combine machine efficiency with human reasoning. They are designed to strengthen cognitive engagement, not reduce it.

The paper concludes with implications for production risk management. Risk assessments rely on contextual understanding, anticipation of behaviour and practical judgement. AI can support elements of the process but cannot replace the depth of thought required for safety-critical decisions. Tools must therefore be designed to keep experts engaged in the final and most important stages of decision-making.

1. Introduction

The rapid adoption of generative AI has prompted widespread debate about whether these tools are reshaping our cognitive habits. The question often appears in media coverage as: Is AI making us stupid? While simplistic, it captures a real concern. If AI increasingly handles planning, analysis and synthesis, what happens to the underlying mental processes we rely on when we think critically?

This question is especially important in creative, judgement-based industries such as TV and film risk management. As a company operating in this domain, Secret Compass has developed its own AI risk-assessment tool. Our position is clear: AI can accelerate work, but it cannot replace human judgement. Understanding why requires a close look at the emerging science.

2. Cognitive impacts of AI use

2.1 MIT Study (2025): Cognitive engagement and cognitive debt

A June 2025 MIT study asked adults to write essays using either ChatGPT, a search engine, or no digital tools (referred to as the brain-only group). Researchers examined electrical brain activity and analysed linguistic patterns. Key findings include significantly lower cognitive engagement among AI users, weaker recall of quotes they had incorporated, and reduced ownership over their work. When participants who had used AI later wrote without it, their cognitive engagement remained low.

The researchers proposed the concept of cognitive debt: prolonged reliance on AI may limit the user's ability to re-engage deeply when required to think independently.

However, the study has limitations. Only 18 participants completed all writing sessions. Other researchers argue the results may reflect a familiarisation effect. The brain-only group completed repeated non-AI tasks and therefore developed familiarity and efficiency. The AI group completed only one non-AI session, limiting direct comparison.

2.2 British Journal of Educational Technology Study (2024): Metacognitive effort and AI dependence

A December 2024 randomised study compared students receiving support from ChatGPT, a human expert, a writing analytics tool, or no tool. Motivation levels were consistent across all groups. However, AI-assisted learners demonstrated different self-regulated learning sequences, showing fewer metacognitive steps such as planning and evaluating. While AI users improved their essay scores in the short term, they showed no significant increase in knowledge retention or transfer.

The authors suggest AI tools may promote metacognitive laziness. Users complete tasks but engage less deeply in the underlying thinking.

2.3 Microsoft Study (2025): Shifting cognitive load

A large Microsoft study in April 2025 examined 936 real-world examples of AI use across 319 knowledge workers. AI increased efficiency, but shifted cognitive load away from exploration and problem-solving. Workers moved from active reasoning to verification. As confidence in AI rose, critical thinking effort decreased. The researchers describe a risk of long-term skill erosion if AI becomes the default problem-solver.

2.4 Microsoft Research Commentary (2024): Narrower ideation and weakening cognitive habits

A 2024 Microsoft commentary synthesising existing research noted that AI can enhance creativity when used deliberately, but more often leads to narrower ideas, less critical evaluation, and reduced retention of information. Workers recall less of what they write or read when AI mediates the process. These effects are particularly evident in low-stakes tasks where users do not feel compelled to check outputs thoroughly.

The commentary argues that AI may help complete tasks faster, but can weaken the cognitive processes that support high-quality thinking.

3. Beyond hallucinations: Understanding workslop

Hallucinations are widely recognised as a limitation of AI tools. A more subtle issue is workslop, a term from Stanford Social Media Lab and BetterUp Labs. Workslop describes AI-generated content that is coherent and polished but lacks depth, nuance or meaningful substance. It creates the appearance of productivity without contributing real insight.

Workslop is particularly challenging because it is plausible. It is easy for users to accept at face value. When AI becomes a crutch, users engage less deeply and reduce the mental steps involved in planning, evaluating and revising. This is a metacognitive problem, not merely a factual one.

In safety-critical fields, shallow or incomplete reasoning can introduce operational risks.

4. Hybrid intelligence: A framework for responsible AI use

4.1 Defining hybrid intelligence

Hybrid intelligence blends machine capabilities such as speed and pattern detection with human capabilities such as judgement, creativity, context and ethical reasoning. Its purpose is not to replace human thinking, but to extend it.

4.2 Cognitive fitness and AI

AI influences cognitive effort. The goal is to use AI intentionally so it reduces cognitive load where appropriate, without replacing the metacognitive work required for reasoning, interpretation and decision-making.

AI can:

- generate ideas users may not have considered
- accelerate administrative tasks
- Identify hazards or patterns
- free cognitive capacity for complex thinking

But only if users remain actively engaged.

4.3 Designing AI with humans in the loop

A responsible hybrid system requires deliberate design choices. These include requiring user input, encouraging revision, providing transparent explanations, and making the user responsible for final decisions. Without such design considerations, users may default to the passive engagement patterns identified in the research.

5. Implications for risk management in TV and film

Risk assessments rely on contextual understanding, anticipation of behaviour, interpretation of creative intent and sensitivity to ethical and practical considerations. These are judgement-led tasks.

AI can support parts of the process but cannot replicate lived experience or the dynamic decision-making required on high-risk shoots. If users disengage cognitively, as suggested by several studies, the outcome is not only a theoretical concern. It may affect safety.

For this reason, AI in risk management must strengthen rather than weaken human thinking. It should accelerate the predictable 80 percent of the work while reinforcing engagement with the remaining 20 percent that carries the real responsibility for safety. This principle is embedded in Secret Compass AI through a hands-on-the-wheel approach.

6. Conclusion

The question of whether AI makes us stupid is simplistic but reflects a legitimate concern. Multiple studies show that AI use can reduce cognitive effort, narrow idea generation and weaken metacognitive habits if used passively. In industries such as TV and film production, where judgement and nuance determine safety outcomes, these risks are significant.

AI must therefore be used intentionally and designed to support hybrid intelligence. When paired with active human oversight, AI can extend thinking, increase efficiency and free cognitive capacity for deeper reasoning. When used passively, it risks weakening the skills that matter most.

AI does not inherently make us less capable. But it can lead us to think less. Responsible design and engaged use are essential to prevent that outcome.

References

MIT (2025)
Kosmyna, N., Hauptmann, E., Yuan, Y.T., Situ, J., Liao, X-H., Beresnitzky, A.V., Braunstein, I. & Maes, P. (2025) *Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task*. arXiv preprint arXiv:2506.08872. Available at: https://arxiv.org/abs/2506.08872

British Journal of Educational Technology Study (2024)

Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X. & Gašević, D. (2024) 'Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance', *British Journal of Educational Technology*, Advance online publication. https://doi.org/10.1111/bjet.13544

Microsoft (2025)
Lee, H-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R. & Wilson, N. (2025) '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 CHI 2025*. Available at: 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/

Microsoft Research Commentary (2024)
Tankelevitch, L., Kewenig, V., Simkute, A., Scott, A.E., Sarkar, A., Sellen, A. & Rintel, S. (2024) 'The Metacognitive Demands and Opportunities of Generative AI', in *Proceedings of CHI '24*. ACM. Available at: https://www.microsoft.com/en-us/research/publication/the-metacognitive-demands-and-opportunities-of-generative-ai/

Stanford Social Media Lab & BetterUp Labs (Workslop concept)
Perry, E. (2025) 'The hidden cost of AI "workslop" — and how leaders can fix it', *BetterUp Business Blog*, 29 September. Available at: https://www.betterup.com/blog/hidden-costs-workslop (accessed [27/11/25]).

Hybrid intelligence supporting literature (optional, verifiable source)
Dellermann, D, Calma, A, Lipusch, N, Weber, T, Weigel, S & Ebel, P 2019, 'The future of human–AI

collaboration: A taxonomy of design knowledge for hybrid intelligence systems', *Proceedings of the 52nd Hawaii International Conference on System Sciences*, pp. 1–10.

Niederhoffer, K., Rosen Kellerman, G., Lee, A., Liebscher, A., Rapuano, K. & Hancock, J.T. (2025) 'AI-Generated "Workslop" Is Destroying Productivity', *Harvard Business Review*, 22 September. Available at: https://hbr.org/2025/09/ai-generated-workslop-is-destroying-productivity (accessed [27/11/25])..

Gerlich, M. (2025) 'AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking', SSRN Working Paper, 3 January. Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5082524 (accessed [27/11/25]).

Walther, C.C. (2025) 'Why Hybrid Intelligence Is the Future of Human-AI Collaboration', *Knowledge at Wharton*, 11 March. Available at: https://knowledge.wharton.upenn.edu/article/why-hybrid-intelligence-is-the-future-of-human-ai-collaboration/ (accessed [021225]).

Notes:

AI studies: In June 2025 MIT published the results of a study where adults were asked to write a series of three essays using either AI (ChatGPT), a search engine, or their own brains ("brain-only" group). The team measured cognitive engagement by examining electrical activity in the brain and through linguistic analysis of the essays. The cognitive engagement of those who used AI was significantly lower than the other two groups. This group also had a harder time recalling quotes from their essays and felt a lower sense of ownership over them. Interestingly, participants switched roles for a final, fourth essay (the brain-only group used AI and vice versa). The AI-to-brain group performed worse and had engagement that was only slightly better than the other group's during their first session, far below the engagement of the brain-only group in their third session. The authors claim this demonstrates how prolonged use of AI led to participants accumulating "cognitive debt." When they finally had the opportunity to use their brains, they were unable to replicate the engagement or perform as well as the other two groups. Cautiously, the authors note that only 18 participants (six per condition) completed the fourth, final session. Therefore, the findings are preliminary and require further testing.

1. Other researchers suggest These results do not necessarily mean that students who used AI
2. accumulated "cognitive debt." In our view, the findings are due to the particular design of the study.
3. The change in neural connectivity of the brain-only group over the first three sessions was likely the result of becoming more familiar with the study task, a phenomenon known as the familiarization effect. As study
4. participants repeat the task, they become more familiar and efficient, and their cognitive strategy adapts accordingly. When the AI group finally got to "use their brains," they were only doing the task once. As a result, they were unable to match the other group's experience. They achieved only slightly better engagement than the brainonly group during the first session. To

fully justify the researchers' claims, the AI-to-brain participants would also need to complete three writing sessions without AI.

In another paper published in the British journal of educational technology in Dec 2024 conducted a randomised experimental study and compared learners' motivations, self-regulated learning processes and learning performances on a writing task among different groups who had support from different agents, that is, ChatGPT (also referred to as the AI group), chat with a human expert, writing analytics tools, and no extra tool. A total of 117 university students were recruited, and their multi-channel learning, performance and motivation data were collected and analysed. The results revealed that: (1) learners who received different learning support showed no difference in post-task intrinsic motivation; (2) there were significant differences in the frequency and sequences of the self-regulated learning processes among groups; (3) ChatGPT group outperformed in the essay score improvement but their knowledge gain and transfer were not significantly different. Our research found that in the absence of differences in motivation, learners with different supports still exhibited different self-regulated learning processes, ultimately leading to differentiated performance. What is particularly noteworthy is that AI technologies such as ChatGPT may promote learners' dependence on technology and potentially trigger “metacognitive laziness”.

We conducted a randomised experimental study in the lab setting and compared learners' motivations, self-regulated learning processes and learning performances among different agent groups (AI, human expert and checklist tools).
We found that AI technologies such as ChatGPT may promote learners' dependence on technology and potentially trigger metacognitive "laziness", which can potentially hinder their ability to self-regulate and engage deeply in learning.
We also found that ChatGPT can significantly improve short-term task performance, but it may not boost intrinsic motivation and knowledge gain and transfer.

Microsoft surveyed 319 knowledge workers who use GenAI tools (e.g.,ChatGPT, Copilot) at work at least once per week, to model how they enact critical thinking when using GenAI tools, and how GenAI affects their perceived effort of thinking critically. Analysing 936 real-world GenAI tool use examples our participants shared, we
find that knowledge workers engage in critical thinking primarily to ensure the quality of their work, e.g. by verifying outputs against external sources. Moreover, while GenAI can improve worker efficiency, it can inhibit critical engagement with work and can potentially lead to long-term overreliance on the tool and diminished skill for independent problem-solving. Higher confidence in GenAI’s ability to perform a task is related to less critical thinking effort.
When using GenAI tools, the effort invested in critical thinking shifts from information gathering to information verification; from problem-solving to AI response integration; and from task execution to task stewardship. Knowledge workers face new challenges in critical thinking as they incorporate GenAI into their knowledge
workflows. To that end, our work suggests that GenAI tools need to be designed to support knowledge workers’ critical thinking by addressing their awareness, motivation, and ability barriers.

In Sept 2024 Microsoft released a thought piece in their research blog - Research in the field has shown mixed outcomes: some studies suggest that, when used in the right way, AI can actually enhance creativity and critical thinking, while others show the opposite. Our own findings highlight the risks. Studies show that using generative AI can lead knowledge workers to produce a narrower range of ideas, put less effort into critical thinking, and retain less of what they write or read. Our survey-based study suggests that when people view a task as low-stakes, they may not review outputs as critically; however, when the stakes are higher, they naturally engage in more thorough evaluation. Even memory – the ability to recall what we've worked on – can be affected when the processes of work are intermediated by AI. In other words, we may get the job done faster, but at the risk of weakening our cognitive "muscles."

Taken together, these emerging findings converge on a core insight: AI can accelerate work, but only by shifting — and often reducing — the cognitive effort humans invest. Without guardrails, this shift risks creating habits of over-reliance, diminishing deep thinking, and weakening the self-regulatory and evaluative skills essential for high-quality creative and analytical work. For industries like TV and film, where risk assessments require contextual judgment, nuance, and ethical sensitivity, these findings underscore the need for AI tools that enhance human thinking rather than replace it — and for users who remain actively engaged in the final, critical 20% of decision-making.

The distinction is subtle but profound. AI as an assistant is about speed and efficiency: finishing tasks more quickly. AI as a tool for thought is about depth: creating outcomes of higher quality, asking more pertinent questions, and exploring unknowns. This matters because knowledge work isn't just about execution; it's about judgment, sense-making, and creativity. If AI intermediates these processes, we risk becoming what Sarkar refers to in his TED talk as "intellectual tourists" – visiting ideas but never truly inhabiting them.

Instead, tools for thought can preserve and even strengthen the habits and capacities of mind that make complex work possible. Tools for thought encourage us to wrestle with ambiguity, reflect on our reasoning, and stay intellectually engaged.