



CLARITY
WORKFLOWS

PERSPECTIVES



PERSPECTIVE 02

Beyond Prompting

Using AI as a collaborative thinking partner.



Christian Adshead
Founder, Clarity Workflows

ABSTRACT

Generative AI is often presented as a faster route to an answer. That framing is useful, but incomplete. Its more significant value may lie in the way it changes the process of thinking itself: making it easier to explore alternatives, expose assumptions, organise complexity and test an idea from several perspectives before acting. This paper argues for a human-led model of AI collaboration in which the technology contributes speed, breadth and intellectual stimulus, while the human retains judgement, accountability and authorship. Drawing on psychological research and practical experience, it considers the difference between cognitive offloading and cognitive augmentation, the risks created by fluent but unreliable outputs, and the importance of productive disagreement. It concludes that the strongest use of AI is not thinking less, but creating the conditions to think better.

KEYWORDS

Human-AI collaboration · Critical thinking · Cognitive offloading · Metacognition · Knowledge work · Creative judgement · Organisational learning

Introduction: beyond better answers

Most people think using AI is about getting better answers.

I do not.

At least, I no longer think that is the most important part of it. Over time, I have come to see AI less as an answer engine and more as a thinking environment: a place in which an incomplete idea can be examined, challenged, reorganised and developed before it becomes a decision, a document or a piece of communication.

The distinction matters. Asking AI to produce something and using AI to think about something may look similar from the outside, but they are different activities. The first can reduce effort. The second can improve judgement – provided the human remains actively involved.

Research already suggests that generative AI can improve speed and output quality on many knowledge tasks. Noy and Zhang found productivity and quality gains in professional writing tasks, while Dell'Acqua and colleagues demonstrated substantial benefits for work that fell within AI's capability frontier. The same research also revealed the opposite effect outside that frontier: people using AI became less likely to reach the correct answer when the system produced persuasive but misleading assistance.^{1,2}

The value of AI is not that it thinks instead of me. It is that, used well, it helps me think better than I would have alone.



Three relationships with AI

AI can be used in at least three distinct ways: as a delegate, as an assistant and as a collaborator.

As a delegate, it is given a task and expected to return a finished product. This is the fastest mode, but it creates the greatest risk of intellectual distance. The user may receive competent-looking work without fully understanding the assumptions, omissions or choices inside it.

As an assistant, AI helps with defined parts of a task: summarising material, proposing alternatives, checking structure, creating first-pass research categories or adapting language for a different audience. The human remains in control of the overall direction, but some cognitive work is deliberately transferred to the tool.

As a collaborator, the relationship becomes iterative. The human brings context, intent, values, experience and accountability. The AI contributes speed, breadth, pattern recognition and the ability to generate multiple lines of enquiry quickly. Neither contribution is sufficient on its own. The useful output emerges through exchange.

My own use increasingly sits in this third category. I rarely accept the first response as the final one. I ask what assumptions have been made. I request a counterargument. I ask the system to separate evidence from inference, identify what is missing, test the reasoning from another stakeholder's perspective and explain where its confidence may exceed the available information.

Fluency is not judgement

Generative AI has a persuasive quality: it can express uncertain or incomplete reasoning in language that feels coherent, calm and authoritative. That fluency is useful, but it can also disguise weakness.

A sentence may be elegant without being accurate. A recommendation may be plausible without being appropriate. A strategy may be internally consistent while resting on the wrong understanding of the customer, the organisation or the problem itself.

This is not a new challenge in human–technology interaction. Research on automation has long distinguished between appropriate use and misuse, including the tendency to rely on automated systems beyond the circumstances in which they are dependable.³

The challenge with generative AI is that the boundary is rarely visible. Dell'Acqua and colleagues describe a 'jagged technological frontier': tasks that appear similar to a human observer can sit on different sides of the model's capability boundary. AI may perform exceptionally well on one and fail convincingly on the next.²

A confident answer is not the same thing as a reliable answer. Human judgement begins where surface fluency stops being enough.



Productive friction

The most valuable AI conversations are not necessarily the smoothest ones. Agreement is comfortable, but it does not always improve thinking. Sometimes the useful moment is the point at which the system challenges my preferred interpretation — or I recognise that its answer is too agreeable, too generic or too detached from reality.

I think of this as productive friction. It is the deliberate introduction of challenge into a process that might otherwise move too quickly towards closure.

In practice, that may involve asking AI to argue against a proposal, identify how a sceptical customer would respond, distinguish emotional appeal from unsupported hype, or show how the same message could be interpreted by people with different priorities. It may also involve rejecting its suggestion entirely.

The purpose is not to make every task longer. It is to spend cognitive effort where it has value. Routine transformation can be accelerated. Consequential judgement should not be surrendered.

Cognitive offloading or cognitive augmentation?

Psychology has long recognised that people use external tools to reduce cognitive demand. Calendars, notes, maps and checklists allow information to be stored outside the mind so that limited attention and working memory can be used elsewhere. Risko and Gilbert describe this process as cognitive offloading.⁴

AI extends that possibility dramatically. It can hold context, compare documents, reorganise information, suggest language and maintain several possible lines of reasoning at once. Used carefully, this can free cognitive capacity for synthesis, judgement and creativity.

But offloading can become abdication. Recent research with knowledge workers found that greater confidence in generative AI was associated with less critical-thinking effort, while greater confidence in one's own ability was associated with more. The researchers also observed that AI changes the nature of critical thinking: away from producing every element directly and towards verification, integration and task stewardship.⁵

That finding captures an important distinction. The question is not whether thinking has been transferred, because tools have always changed where thinking happens. The question is whether the human remains cognitively present — understanding the task, evaluating the output and taking responsibility for the result.



Why this matters for copy and content

Copy is not simply the arrangement of persuasive words. Effective communication begins with an interpretation of people: what they want, what they fear, what they already believe, what they resist and what would make the next step feel credible.

AI is exceptionally useful for widening that enquiry. It can generate alternative customer hypotheses, organise review themes, compare positioning, test hooks, explore objections and adapt a message across formats. It can help a writer move rapidly between the perspectives of customer, strategist, sceptic, editor and brand guardian. It can also flatten the work. Research by Doshi and Hauser found that generative AI increased the judged creativity of individual stories while making the collection of stories more similar to one another. In other words, individual performance improved, but collective variety narrowed.⁶

That is a significant warning for brands. When the same systems are used in the same way, competent output can converge towards a polished average. Taste, specificity and original observation become more important, not less.

My preferred approach is therefore to use AI before, around and after the writing — but not to let it erase the human source of the work. I use it to investigate the emotional tension behind a decision, identify competing interpretations, test whether a claim has earned its confidence and refine the route from attention to action. The final language must still sound as though it belongs to a real brand speaking to a real person.

A practical collaboration loop

My own workflow is not a fixed prompting formula. It is a repeatable reasoning loop:

1. **Frame the problem.** Define the decision, audience, desired response, constraints and available evidence before asking for output.
2. **Expand the field.** Use AI to generate interpretations, angles, objections, structures and possibilities that might otherwise be missed.
3. **Introduce challenge.** Ask what is weak, assumed, generic, unsupported or likely to be resisted.
4. **Verify and ground.** Check material claims against source evidence, lived context, customer language or operational reality.
5. **Exercise judgement.** Select, reject, rewrite and take responsibility for the final direction.
6. **Review the effect.** Assess not only whether the output reads well, but whether it creates the intended understanding, feeling or action.

This loop keeps the technology active without making the human passive. It also recognises that good work is rarely produced by a single instruction. It develops through clarification, tension, selection and revision.



Implications for organisations

Organisations adopting AI often begin with access: which tool should be purchased, who should receive a licence and which tasks could be completed faster. Those questions matter, but they are not the starting point.

The more useful starting point is to understand the work itself. Where is judgement essential? Where does staff knowledge remain undocumented? Which tasks consume cognitive effort without adding value? Where could AI broaden thinking, and where might it create false confidence?

Human–AI interaction research consistently emphasises the need to make capabilities, limitations, correction and control understandable to users. The objective is not frictionless automation at any cost, but a relationship in which people can anticipate, evaluate and appropriately steer the system.⁷

The organisations that benefit most will not necessarily be those that automate the greatest number of tasks. They will be those that design the clearest division of labour between technological capability and human responsibility.

Conclusion

AI can make work faster. That is already valuable.

But speed is not its most interesting contribution.

Used as a serious thinking partner, AI can help people examine a problem from more angles, recognise assumptions earlier, organise complexity, test language and reach a more considered judgement. It can create space for thought — but only if the user resists the temptation to disappear from the process.

The central skill is therefore not prompting in isolation. It is stewardship: knowing what to delegate, what to challenge, what to verify and what must remain distinctly human.

The future of effective AI use will not be defined by people who ask machines to think for them. It will be shaped by people who learn how to think with them.

*For me, that is what collaboration means.
Not replacing thought.
Not performing it theatrically after the system has
done the work.
Thinking — together.*



References

1. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
2. Dell'Acqua, F., McFowland, E., III, Mollick, E., Lifshitz, H., Kellogg, K. C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2026). Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science*, 37(2), 403–423. <https://doi.org/10.1287/orsc.2025.21838>
3. Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2), 230–253. <https://doi.org/10.1518/001872097778543886>
4. 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>
5. 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 the 2025 CHI Conference on Human Factors in Computing Systems* (Article 1121, pp. 1–22). Association for Computing Machinery. <https://doi.org/10.1145/3706598.3713778>
6. Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), eadn5290. <https://doi.org/10.1126/sciadv.adn5290>
7. Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Article 3, pp. 1–13). Association for Computing Machinery. <https://doi.org/10.1145/3290605.3300233>



Christian Adshead

FOUNDER, CLARITY WORKFLOWS

Christian Adshead brings experience across psychology, digital media, service development, operational leadership, project work and organisational improvement. His work focuses on making complex information, processes and decisions clearer, more practical and easier to act upon.

Publication note. Human-authored, AI-assisted: AI helped shape structure and presentation; final content and references were reviewed by the author.

CLARITY WORKFLOWS PERSPECTIVES — PERSPECTIVE 02 OF AN ONGOING SERIES

If this paper has raised questions relevant to your organisation, the Clarity Workflows Business Check offers a complimentary structured conversation to explore whether AI could usefully support your team's thinking — and where it should not.

To continue the conversation, please contact christian@clarityworkflows.com or visit clarityworkflows.com

