



PERSPECTIVE 03

Before the Prompt

Where originality, experience and human judgement enter AI collaboration.



Christian Adshead
Founder, Clarity Workflows

ABSTRACT

Generative AI can produce polished language, unexpected connections and genuinely useful ideas. The challenge is not proving that humans are creative and machines are not. It is deciding where human judgement, experience and responsibility should remain when the machine becomes an increasingly capable collaborator.

KEYWORDS

Human-AI collaboration · Originality · Creative judgement · Convergence · Cognitive offloading · Social conformity · Knowledge work

Introduction

There is a growing conversation about AI-assisted content beginning to converge: the same structures, the same phrases, the same reassuring conclusions. Scroll through LinkedIn for long enough and it is not difficult to see why the concern exists.

But there is a danger in making the explanation too simple. The problem is not merely that AI is incapable of originality and therefore everything it touches becomes generic. That is comforting rhetoric, but it is not quite true.

AI can generate ideas that people judge to be creative. It can make associations we have not made, suggest directions we had not considered and improve the quality of an individual piece of work. Research published in *Science Advances* in 2024 found that access to generative-AI ideas increased ratings of creativity in short stories, particularly for less-creative writers. The same study also found that the stories became more similar to one another. That tension matters.^[1]

So perhaps the better question is not, 'Can AI be creative?' It is: 'What happens to human originality when we collaborate with something that can generate plausible ideas almost instantly?'

Curiosity, Not Commands

In the previous Clarity Workflows Perspective, *Beyond Prompting*, I argued that the more useful relationship with AI is not simply command and response. It is collaboration: questioning, testing, revising, challenging and developing an idea over time. But collaboration raises another question. If the AI is doing more than formatting or polishing our work, what should we still bring to the process ourselves?

One answer is deceptively simple: something. An observation. A frustration. A contradiction. A half-formed idea. An experience. A question you cannot quite answer.

Compare these two starting points:

ASK AI TO PROVIDE THE THINKING

'Give me ten original LinkedIn posts about AI.'

BRING YOUR OWN THINKING FIRST

'I keep noticing people criticising AI-generated content for sounding the same, but I wonder whether humans were already converging around safe ideas before AI arrived. What am I missing?'

The second prompt is not necessarily more sophisticated technically. It is more human. It brings a thought into the conversation and asks AI to work on it rather than asking AI to provide a thought to adopt.

Point of View?

This distinction matters because the useful human role is not simply 'being more intelligent' or 'being more creative' than AI. That framing may become increasingly difficult to sustain as the technology develops.

AI can produce a persuasive argument for one side of a question and, seconds later, an equally persuasive argument for the other. It can generate unconventional possibilities. It can imitate styles, combine concepts and help create something that neither the user nor the model would have reached in exactly the same way alone.

What it cannot do on our behalf is decide which conclusion we are prepared to own.

That requires judgement. Do I actually believe this? Have I understood it? What assumptions am I making? Am I accepting this answer because it is convincing, or because it conveniently confirms what I already wanted to say? Would I defend this position if someone challenged it?

Those questions become more important, not less, as AI becomes more capable.

The Human Condition

There is another reason to be uncomfortable simply blaming AI for generic content.

Humans converge too.

Long before generative AI, we watched what other people were doing. We noticed which ideas gained approval, which behaviours were rewarded, which language signalled credibility and which opinions were likely to create discomfort. We learned how people in our professions spoke, what successful businesses looked like, what good leadership was supposed to sound like and, increasingly, what performed well online.

There is nothing particularly sinister about this. It is part of being human.

Social psychology has long distinguished between different forms of social influence. Sometimes we look to other people because we are uncertain and assume that the group may know something we do not. Sometimes we conform more publicly because belonging matters and social rejection carries a cost. And sometimes the knowledge that our ideas will be evaluated changes what we are prepared to offer in the first place.^[2]

These are not simply signs of weak thinking. They are deeply human responses to uncertainty, belonging, status and risk.

But our current environment gives those responses unusually powerful feedback.

We can see what gets liked. What gets shared. What receives praise. What attracts criticism. What gets promoted. What sells. What other organisations are doing. Which phrases are circulating. Which subjects are suddenly everywhere.

And when value becomes closely associated with visible output — more content, more engagement, more productivity, more measurable results — it becomes understandable that we gravitate towards things that already appear to work.



That creates an interesting tension around originality.

Someone may be perfectly capable of thinking differently and still hesitate to say something different. Not necessarily because they lack courage, but because the unfamiliar carries uncertainty. A genuinely undeveloped idea has no engagement data attached to it. Nobody has validated it yet. There is no guarantee that other people will understand it, agree with it or even consider it worthwhile.

The familiar is safer.

AI did not create that human condition. But it can make responding to it remarkably efficient.

If a particular conversation is already gaining traction, I can ask an AI system to help me contribute to it. It can identify the established arguments, produce a credible structure, adopt the vocabulary surrounding the subject and generate something competent in seconds. The technology is doing exactly what I have asked it to do.

Perhaps the more uncomfortable question is why I asked.

There are increasingly familiar posts pointing out the familiar structures, expressions and arguments appearing in AI-assisted writing. That criticism is not pointless — it raises an important problem.

But there is also a paradox. The criticism itself can begin to converge. The same observations appear. Similar examples are used. A recognisable conversation develops around the problem of recognisable content.

That does not make the criticism wrong. Quite the opposite — those conversations helped lead me to the question I am exploring here.

But perhaps convergence is not simply evidence that AI is making us less original. Perhaps AI is entering a social system in which human beings were already highly responsive to what other human beings were thinking, saying and rewarding.

If so, the challenge is not merely to make AI produce more original content. It is to remain sufficiently present in the process to notice when we are choosing familiarity for ourselves.

From A to B — Fast!

There may be another distinctly human factor in all of this: our relationship with outcomes, and with what those outcomes appear to say about us.

We often measure success through what we achieve and, consciously or otherwise, through how those achievements are perceived by other people. Was the work good enough? Did the client like it? Did the post perform? Did I look competent? Did I get the result?

That is not something AI brought into the room. It is part of the human condition. We care about belonging, reputation, approval, security and status. We can worry about being rejected, criticised, overlooked or simply thought less of.



A large language model does not have those social anxieties. It is not sitting behind the screen thinking that it had better produce something original or I may stop liking it. It does not fear being ostracised, losing its job or looking foolish in front of its peers. Nor is it secretly trying to make everybody sound the same.

It generates language from patterns learned across enormous amounts of human material and from the context we give it. When the context is broad and the outcome is conventional, a plausible, coherent and broadly applicable answer is often exactly what we have asked it to produce.

And humans have always looked for faster ways to get from A to B. We automate. We standardise. We create templates. We invent tools. We remove unnecessary steps. Much of human progress depends on doing precisely that.

Generative AI is extraordinarily good at shortening the distance between an intention and an output.

The Missing Material

Suppose I have met a prospective client and need to send a follow-up email. I could tell an AI system who the client is, what their business does and what I want the email to achieve. Seconds later, I have a polished professional message. The outcome has been achieved.

But perhaps something has disappeared on the journey.

The AI does not know that the client bought me a coffee. It does not know that I particularly enjoyed one part of our conversation. It did not notice their bright red shoes, hear the story that made me laugh or experience the moment when the conversation moved from a business transaction into something more human.

Unless I bring those details into the collaboration, the finished email cannot contain them. It may be professionally competent and completely appropriate. It may also be interchangeable with a thousand other professionally competent and completely appropriate emails.

That is not because the AI has failed. It has fulfilled the brief. Perhaps the problem is that the brief contained the outcome but very little of the experience that produced it.



A great deal of human originality does not begin at a desk, neatly labelled as an idea. It arrives while people are walking, running, showering, driving, talking, failing, arguing, cooking or lying awake thinking about something that happened three days earlier.

Experience, conversation, memory, failure, curiosity, boredom, observation and unexpected association are not sentimental extras around the creative process. They are the raw material from which ideas are born.

AI can participate once we bring that material into the conversation. It can question it, connect it, challenge our interpretation and help us articulate something we have not yet fully formed.

But if our first interaction with AI begins only with the destination — I need a post, I need an email, I need a presentation, I need five ideas — perhaps we should not be surprised when the quickest route leads somewhere familiar.

Originality in an AI-assisted world therefore begins before the prompt. It begins in what we notice, what we care about, what we have experienced and what we are willing to bring with us when we start the conversation.

The danger is not that AI removes the human journey. It is that our focus on the destination encourages us to remove ourselves from it.

Put the Friction Back

Much of the productivity conversation around AI is understandably about reducing friction. Draft faster. Summarise faster. Research faster. Generate options faster.

But not all friction is waste.

Sometimes the uncomfortable period before we know what we think is part of thinking. The blank page can be frustrating precisely because an idea has not yet settled. If our first response to that discomfort is always to ask AI to fill the space, we may remove the part of the process in which our own judgement would otherwise have developed.

A better collaboration deliberately introduces useful friction. Instead of asking AI to tell us what to think, ask it to question what we think:

QUESTIONS WORTH ASKING AI

- ? What is weak in this argument?
- ? What assumption am I making?
- ? What would somebody who disagrees with me say?
- ? Is this observation genuinely distinctive, or am I repeating a fashionable argument?
- ? What evidence would change my mind?
- ? What have I not considered?

Recent experimental research supports this distinction. A large preregistered study presented at CHI 2026 compared model-led and human-led modes of AI co-creation. When the model proactively refined participants' ideas, quality improved but diversity and perceived ownership fell. In a human-led question mode, ideas remained high quality while diversity increased.^[3]

Active human involvement is not a sentimental add-on to AI collaboration. It changes the outcome.^[4]

Something Is Better Than Nothing

This suggests a different approach. Bring something human to the starting point – something noticed, experienced, questioned or curious. Then work with AI from there:

- 1 Ask questions before asking for answers.** Let AI interrogate the idea, expose gaps and show you perspectives you have not considered.
- 2 Widen the field rather than settle the conclusion.** AI is exceptionally good at generating alternatives. That does not mean one of those alternatives has to become yours.
- 3 Disagree with it.** Reject suggestions. Go backwards. Change direction. If every response is immediately accepted, it is difficult to call it collaboration.
- 4 Keep the final judgement human.** The important question is not simply whether AI helped produce the words. It is whether you understand, believe and are willing to take responsibility for what you say.

Original, Not Groundbreaking

There is another trap here: believing that originality requires an idea nobody else has ever had. That is an impossibly high standard for most work. Human thought has always been cumulative. We read, listen, remember, borrow, challenge, combine and reinterpret. The originality can lie in the particular route through those influences – the connections made, the experiences brought to them, the questions asked and the judgement applied.

It is a lower bar than we sometimes assume. And a more honest one.



The Things We Think and Do Not Say

Thinking about this brought to mind, perhaps unexpectedly, the 1996 film *Jerry Maguire*, starring Tom Cruise and Renée Zellweger.

Early in the film, Jerry experiences what he describes as a moment of clarity. He stays up through the night writing a mission statement — *The Things We Think and Do Not Say* — challenging some of the assumptions underpinning the sports-agency business around him. Then he prints it. And puts a copy in everyone's pigeonhole.

That part matters. Having the thought privately carries relatively little risk. Writing it down begins to make it real. Putting it into the hands of the people whose system it challenges makes him accountable to it. The immediate outcome is not admiration for his originality. He loses his job.

It is fiction, and the eventual Hollywood outcome is rather more satisfying, but the behavioural tension is very real: there can be a cost attached to departing from the assumptions of the group.

Business history provides a less theatrical example. In 1975, Kodak engineer Steven Sasson developed the world's first self-contained digital camera. The significance of that story is not simply that Sasson had an original idea. The original thinking was already sitting inside one of the world's dominant photography companies.

The more interesting question is what that idea implied. Kodak's enormously successful business was built around film. Digital photography did not simply offer the company another product. It challenged some of the assumptions on which the existing model depended. Sasson continued developing digital camera technology, and years later Kodak declined to commercialise a self-contained digital SLR developed by Sasson and a colleague, with protection of film sales forming part of the decision.^[5]

The problem, therefore, was not an absence of creativity. An innovative idea existed. What was difficult was what accepting that idea might require the existing system to reconsider.

Not every unconventional idea is a good one. Organisations cannot sensibly abandon their business models every time somebody proposes something different. Equally, disagreeing with the majority does not magically make an idea original, intelligent or correct.

But genuinely different thinking often contains uncertainty because, by definition, there is less existing evidence telling us how other people will respond to it.

That applies at a far smaller scale when we publish something online. Most of us are not inventing digital photography or writing mission statements that threaten entire industries. But we are still making small decisions about whether to follow an established conversation or introduce something less certain into it.



Do I write about the subject everybody is currently discussing? Do I use the formulation I already know people recognise? Do I ask AI what is trending and generate my contribution from there?

Or do I begin with the odd observation I made yesterday, the question I cannot quite answer, the thing I disagree with, or the thought I am not yet sure is even right?

This is where the way we collaborate with AI starts to matter. Because if both the human and the technology are being asked to move towards what is already familiar, convergence should hardly surprise us.

The alternative is not to reject AI. It is to use the collaboration differently. Bring the uncertain thought. Then ask AI to challenge it. Ask what is weak about it. Ask whether the conventional view may actually be right. Ask what evidence contradicts it. Ask what assumption you have smuggled into the argument.

AI can remove friction extraordinarily well. Perhaps one of the more valuable ways to collaborate with it is occasionally to ask it to put some friction back.

Who Wrote This?

A final disclosure feels appropriate.

AI drafted much of the prose in this Perspective. But the argument did not begin with a prompt asking for an article about AI and originality. It developed through a conversation: an observation about converging content, a challenge to the claim that AI cannot be creative, a question about whether the argument might itself be generic, a detour into conformity and social judgement, a further question about our preoccupation with outcomes and what gets lost when we rush from intention to output, disagreement about where the article was going, and eventually a return to the original question — how should we collaborate with AI without outsourcing the thinking?

The AI contributed information, structure, counterarguments and language. I questioned it, redirected it, rejected parts of it and decided what I was prepared to keep.

That does not prove the finished Perspective is original. In a world of billions of people and increasingly capable generative systems, absolute originality is a difficult claim to make. But perhaps that is not the most useful test. A better test may be:

Did I remain actively involved in the thinking?

Do I understand what I am publishing?

Do I believe it?

Have I challenged it?

Am I prepared to put my name to it?

If the answer is yes, AI has not removed the human from the work. It has participated in the conversation.

*AI can contribute to the conversation.
It should not remove you from it.*

References and Further Reading

- [1] Doshi, A. R. & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28). doi.org/10.1126/sciadv.adn5290
↑ [Back to citation](#)
- [2] Sowden, S., Koletsis, S., Lymberopoulos, E., Militaru, E., Catmur, C. & Bird, G. (2018). Quantifying compliance and acceptance through public and private social conformity. *Consciousness and Cognition*, 65, 359–367. doi.org/10.1016/j.concog.2018.08.009
↑ [Back to citation](#)
- [3] Maier, S. et al. (CHI 2026). Partnering with Generative AI: Experimental Evaluation of Human-Led and Model-Led Interaction in Human-AI Co-Creation. arXiv: 2510.23324
↑ [Back to citation](#)
- [4] Dong, M. & Yakura, H. (2026). Human diversity fuels collective creativity that large language models cannot simulate or sustain. arXiv: 2607.26899
↑ [Back to citation](#)
- [5] National Inventors Hall of Fame. Steven Sasson — inventor profile and legacy account documenting the first self-contained digital camera (1975) and the later self-contained DSLR developed with Robert Hills (1989).
↑ [Back to citation](#)



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.

CLARITY WORKFLOWS PERSPECTIVES – PERSPECTIVE 03 OF AN ONGOING SERIES

Clarity Workflows Perspectives offers short-form thinking on human-led, AI-assisted work, workflow design and the changing relationship between people and technology.

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

