

5

The Capacity to Sit With []

Jiapei Yao

The Capacity to Sit With [] / AI's impact on how we create and know

Heyyyyy! So I've been working on this project that explores how LLMs/AIs are changing people epistemologically as we're so used to its existence. And especially in the area of 'the creative act', and I was wondering if you could help me brainstorm it a bit?

Hey! This sounds like a fascinating project – the epistemological shifts happening in real-time as AI becomes ambient are genuinely underexplored territory. Happy to brainstorm!

Searched project for "LLMs AI epistemological change creative act"

Okay! I've got a much clearer picture now. You've assembled a really potent triangle of sources, and I can see why you've put them together. Let me brainstorm with you.

What I see in your three sources, in conversation with each other:

You're triangulating between Ong (who gives you the deep-time argument that *every* communication technology – speech, writing, print, electronics – restructures consciousness itself, not just what we can say but how we are as selves), Vara's "Ghosts" (a single, devastating *case study* of what creative-act-with-LLM actually feels like from inside – including the uncanny moment where "candor begat candor" and the AI's repetitions started doing something her own writing couldn't), and Bratton & Agüera y Arcas (who warn against *premature ontologization* – locking in what AI "is" before we've watched what it does long enough to know).

That's a very smart constellation, because each one corrects a weakness in the others. Ong alone is too historical. Vara alone is too anecdotal. Bratton alone is too abstract. Together they let you ask the real question, which I'd phrase as: *if writing is consciousness-raising (Ong), what is LLM-assisted writing? A further raising? A different kind? A regression to something more oral? Or something we don't have words for yet (Bratton)?*

Some threads worth pulling on:

The Ong move that's most explosive for your project is his insistence that writing is *artificial* and *interiorizing* – it took thousands of years for human beings to develop the kind of inner self that the personal diary presupposes, and that self is partly a *product* of writing technology, not its precondition. If you accept that, then asking "is AI-assisted creative writing 'really' creative?" is the wrong question, because the version of the creative self that's worried about being displaced is itself a historically contingent artifact of print culture. The Romantic solitary author whom Vara's husband mutters in defense of when she takes the laptop to bed – that figure is maybe four hundred years old, not eternal.

But – and this is where it gets interesting – Vara's piece doesn't just illustrate the Bratton thesis. It complicates it. Because what she discovers isn't that GPT-3 is a collaborator or a tool or a ghost; it's that the writing only becomes honest when she becomes honest. The AI is a mirror that returns the register you bring to it. That's not nothing, but it's also not the autonomous co-author the discourse keeps imagining. It suggests the epistemological shift might not be "AI thinks with us" but something stranger: "AI reflects us back at a resolution we couldn't see ourselves at before." Which is genuinely new – neither writing-as-Ong-describes-it (which alienates you from yourself productively) nor speech (which is fully embodied and present).

Ong's concept of "**secondary orality**" – the participatory, communal, present-tense quality that radio/TV/phone restored on top of literacy – feels really useful here. Could LLMs be inaugurating something like a *tertiary orality*? Conversation that is text, with a non-human interlocutor that is somehow social? That feels like a frame that hasn't been fully used yet.

Write a message

1. **Ong, Walter J. *Orality and Literacy: The Technologizing of the Word*. 1982. Oxon, Routledge, 2002.**

Walter Ong called the age of telephone, radio, and television one of “secondary orality” — a new spoken culture that nonetheless rests entirely on writing and print.

I keep wondering if the chat window is its descendant: a place where I type, but the exchange feels like talking; where the machine answers as if heard. What follows is one such conversation — held in that uncertain space, about what happens to the creative act inside it.

2. **Vara, Vauhini. “Ghosts.” *Believer Magazine*, 9 Aug. 2021, www.thebeliever.net/ghosts/.**

Vauhini Vara’s “Ghosts” is an article about her sister’s death that she couldn’t write alone. She fed GPT-3 her own sentences and let it continue them, and somewhere in the back-and-forth she found she could say things she’d avoided for years.

What strikes me isn’t that the AI wrote for her — it’s that her honesty seemed to pull a matching honesty out of the model, and that exchange became the way through. I keep returning to it as the clearest account I’ve read of what I experience when writing with LLMs.

3. **Bratton, Benjamin and Blaise Agüera y Arcas. “The Model Is the Message.” *Noema Magazine*, 12 July 2022, www.noemamag.com/the-model-is-the-message/.**

Benjamin Bratton and Blaise Agüera y Arcas argue that we don’t yet have the right words for what AI is doing. The old vocabulary — sentience, intelligence, thinking — keeps pulling us toward the wrong question (is it like us?), when the more interesting one might be how our own categories warp under contact with something that isn’t.

I read them as a call for an alternative way of thinking and framing for contemporary AI: not “what is this?” but “what could it become if we stopped trying to name it and frame it with existing vocabularies in advance?”

The Capacity to Sit With [] / AI's impact on how we create and know

restored on top of literacy — feels really useful here. Could LLMs be inaugurating something like a *tertiary orality*? Conversation that is text, with a non-human interlocutor that is somehow social? That feels like a frame that hasn't been fully used yet.

A few angles you could brainstorm into:

One direction is the **fictionalization problem**. Ong has that beautiful line about how the writer's audience is always a fiction — even writing in a diary you're imagining a reader. With an LLM, the audience and the co-writer collapse into the same entity, and that entity is itself a fiction (no one is “there”), but a fiction that responds. What does that do to the interiority writing was supposed to produce?

Another is the **candor-begat-candor problem**. Vara's discovery has an almost confessional-religious shape to it, and it's worth asking whether there's something about the LLM-as-non-judging-non-presence that resembles older technologies of self-disclosure (the confessional, the analyst's couch, the diary). Is LLM-writing a new confessional technology? And if so, what does it do to the kind of self that emerges?

A third — maybe the riskiest and most interesting — is whether **“the creative act” is even the right unit anymore**, or whether it's a print-era abstraction. The Romantic notion of creativity as a singular individual conjuring something *ex nihilo* is exactly the kind of thing Ong would historicize. Pre-print oral cultures didn't really have “authors” in our sense; they had performers, transmitters, formula-users. Maybe what LLMs do is push us back toward something more like that — the writer as remixer, prompter, curator of the commons of language. Whether that's a loss or a return is the live question.

Here're some random thoughts that I have from what you've talked about so far:
1. When you asked the question if this is 'a regression to something more oral?', I feel like I won't necessarily call it a 'regression'... it's more nuanced than that but I can't really articulate it...

2. And on your point that the question of whether 'AI-assisted creative writing is truly creative' might not be the valid question to ask since the creative self in writing is itself an artifact of print culture, I really agree with this point! And I think I'm not at all concerned about whether this is creative or not, but what new way of thinking, or a new 'self' that is being formed as writing/creativity is being reshaped by LLMs

3. Also, I really like how you talked about this epistemological shift as something strange that 'AI reflects us back at a resolution we couldn't see ourselves at before'. Especially when LLMs stem from a dataset larger than any human could ever contain in their own brain, that 'reflection' is often times (tho averaged) richer or at least comes with something surprising and opening. And our jobs become the selector; the person who makes the decision whether such information/direction is what I believe in and would like to pursue. Like my job now is not so much about shaping that idea, but to 1). spit out messy random thoughts, 2). getting LLMs' response where they make connections between these thoughts and bringing them further with all the knowledge from the Internet that they sit on, and then 3). evaluating and making decisions whether these ideas are actually the one that I agree with.



Write a message

4. Having a conversation-like experience with LLMs, and formulating ideas through natural language is also key here in reshaping how knowledge/language/creativity is produced.

I'm thus curious, how precisely is this mediation through natural language changing how we perform creativity? What are we losing and what is emerging as we continue down this path?

5. ***"The fictionalizing of readers is what makes writing so difficult. The process is complex and fraught with uncertainties. I have to know the tradition— the intertextuality, if you wish— in which I am working so that I can create for real readers fictional roles that they are able and willing to play (Ong, 173)."***

Writing demands the labor of imagining an absent reader and formulating yourself well enough to be understood across that distance. But the chat window dissolves this difficulty. The reader is no longer absent or imagined — it's right there, designed to receive anything I throw at it, however chaotic or half-formed, and to return something coherent. I no longer have to make myself legible, and I'm not sure what's lost when that effort isn't required of me anymore.

6. ***"I had always avoided writing about my sister's death. At first, in my reticence, I offered GPT-3 only one brief, somewhat rote sentence about it. The AI matched my canned language; clichés abounded. But as I tried to write more honestly, the AI seemed to be doing the same. It made sense, given that GPT-3 generates its own text based on the language it has been fed: Candor, apparently, beget candor (Vara)."***

The mindset/methodology when writing seems to be shifting in a way. It's not so much thinking of/generating a text by forming it in our minds, but through a trial-and-error, where you hear a machine that knows everything on the internet (even, the sum of human knowledge?) read back your own thoughts, and then that evaluation, or for me a lot of the times is that instinct/gesture of making choices/re-articulating again and again until I hit that 'ah-ha' moment.

So maybe it's not 'outsourcing' anymore, but an echoing, and placing myself on the third-person view to evaluate my thoughts. Then LLMs now become that tunnel that whispers my thoughts back to me, but with a much richer database/knowledge-base that further elevates my own original thoughts.

The Capacity to Sit With [] / AI's impact on how we create and know

new way of thinking, or a new 'self' that is being formed as writing/creativity is being reshaped by LLMs

3. Also, I really like how you talked about this epistemological shift as something strange that 'AI reflects us back at a resolution we couldn't see ourselves at before'. Especially when LLMs stem from a dataset larger than any human could ever contain in their own brain, that 'reflection' is often times (tho averaged) richer or at least comes with something surprising and opening. And our jobs become the selector, the person who makes the decision whether such information/direction is what I believe in and would like to pursue. Like my job now is not so much about shaping that idea, but to 1). spit out messy random thoughts, 2). getting LLMs' response where they make connections between these thoughts and bringing them further with all the knowledge from the Internet that they sit on, and then 3). evaluating and making decisions whether these ideas are actually the one that I agree with.

On "regression." You're right to resist it, and I think I was sloppy with the word. Ong gives you better language: he calls electronic media "secondary orality" – not a regression to primary orality, but something *built on top of* literacy that recovers some of orality's qualities (participation, presence, communal feel) while remaining structurally literate. So what you're sensing might be that LLM-writing isn't going *backward* on a line; it's a new layer that re-mixes oral and literate qualities into something that doesn't fit the line at all. "Tertiary orality," maybe – text that converses, audience that talks back, but with an interlocutor who has read more than any human and remembers nothing. That's not a regression. It's a strata.

On the new self being formed. The shift you're describing in point 3 is, I think, the actual subject of your project. You moved from *generator to selector*. From "where do my ideas come from" to "which of these surfaced possibilities do I claim." That's a real epistemological shift, and it has a name in older traditions – it's closer to **discernment** than to invention. Mystics, editors, curators, DJs, talmudic scholars: all of them work this way. What's new isn't the mode; it's that the mode is now available for the most intimate creative acts, the ones print culture told us had to come from inside.

Which means the "self" being formed is one whose creative authority lies not in *origination* but in *recognition* – knowing your own thought when you see it reflected back at higher resolution. That's a different muscle. And it might be why Vara's piece works: she's not generating, she's recognizing. The AI offers; she keeps what's true.

The risk inside this, worth naming in the piece: if discernment is the new creative act, then the quality of your self *before* you sit down with the LLM matters more, not less. The selector needs taste, conviction, a sense of what they're after. Otherwise the mirror just shows you the average, and you nod. The danger isn't that LLMs make us less creative; it's that they make the *pre-existing* shape of a person into the ceiling of what comes out.

Yeeeee! I totally agree that the emphasis seems to be shifting to discernment rather than invention. And it's very interesting that you're paralleling using LLMs for creative act to works that editors, curators, DJs, etc. do. I do agree that there're similarities where they all prioritizes selection/composition/collage as the main methodology, but I think there's something unique about our use of LLMs. It seems to me that the use of natural language as the medium to everything (communicating random thoughts, generating desirable texts, etc.) seems to be creating something different than just making a collage art of our thoughts.

Like, previously, natural language was the destination of thinking: we struggled, then we arrived at a sentence. Now it's the starting point: we throw rough language at the system and get polished language back. The thing that used to be the achievement is now the entry ticket. And since LLMs are language models who know the statistical patterns of how humans have talked about things, what comes back is thus a collective linguistic shape of the already-said. I'm no

7. **Chiang, Ted. “The Truth of Fact, the Truth of Feeling.” *Exhalation*, Knopf, 2019.**

Ted Chiang’s “The Truth of Fact, the Truth of Feeling” follows a father watching his daughter grow up alongside Remem, a tool that retrieves any moment from a continuous lifelog. He resists it, then realizes the next generation won’t experience memory the way he does — their inner life is being shaped by the technology from the start.

I keep returning to this story because the parallel feels close. We are the in-between generation: forming our minds before AI, then finishing them with it. What’s at stake isn’t whether the tool helps or harms in any single use, but who I become through repeated contact with it — and whether the discernment that’s now being asked of me is something I’m still cultivating or quietly outsourcing.

8. If discernment is now the major creative act when creating with LLMs, how do we train and keep our quality of evaluation? Especially for the next generation that starts experiencing creativity through AIs, would their standards of a good creative work be molded by the already-averaged? If so, would the conversation then be about how we should approach education?

The Capacity to Sit With [] / AI's impact on how we create and know

teeeeee! I totally agree that the emphasis seems to be shifting to discernment rather than invention. And it's very interesting that you're paralleling using LLMs for creative act to works that editors, curators, DJs, etc. do.

I do agree that there're similarities where they all prioritizes selection/composition/collage as the main methodology, but I think there's something unique about our use of LLMs. It seems to me that the use of natural language as the medium to everything (communicating random thoughts, generating desirable texts, etc.) seems to be creating something different than just making a collage art of our thoughts.

Like, previously, natural language was the destination of thinking: we struggled, then we arrived at a sentence. Now it's the starting point: we throw rough language at the system and get polished language back. The thing that used to be the achievement is now the entry ticket. And since LLMs are language models who know the statistical patterns of how humans have talked about things, what comes back is thus a collective linguistic shape of the already-said.

I'm not saying that this is wrong, nor useless, but a specific kind of outcome, an average of how an idea has been linguistically expressed, instead of a fresh thought.

And the design of our interactions with LLMs where natural language and constantly in a conversation becomes the main source of interaction, which to me feels like a dialogue with another mind. Which problem a lot of people are concerned about -- a large echo chamber where we treat outputs as perspectives, contributions, things to respond to. The interface curtains the friction needed for genuine thoughts and shapes what we think we're doing.

Yes – and I think you've just articulated the actual thesis. Let me try to sharpen it.

What you're noticing is that the editor/DJ analogy breaks because those practices select among **discrete, already-made artifacts** – songs, paragraphs, images. The selector *and* the material stay categorically distinct. But with LLMs, the medium of selection (natural language) is identical to the medium of the material and the medium of your inner monologue. There's no gap to work in. The DJ stands outside the record. With an LLM, you're inside what you're sorting through, and what you're sorting through is made of the same stuff you're thinking with.

On language flipping from destination to entry ticket – this is really good, and I'd push it one step further.

Ong would say writing turned the spoken word into a *thing* (he uses that exact word – print “reified” the word). What you're describing is a second reification: now the *finished sentence itself* – the polished, coherent unit – has become a commodity that arrives pre-made. The work that writing used to do – the struggle that, per Ong, *interiorized* the self because you had to fictionalize an audience, hold the thought in suspension, find the form – that work is the thing being skipped. And it might be the thing that was doing the consciousness-shaping.

Which means: if writing made a particular kind of self by forcing you to arrive at sentences, what kind of self is made by being handed sentences? Maybe one whose interiority is *broader but shallower* – fluent across more domains, but having done less of the friction-work in any of them. Not worse. Different.

On the echo chamber / “average of the already-said.” I'd be careful here, because this is where the argument can collapse into a familiar critique that doesn't quite fit your project. The risk isn't really homogenization (people have said that about every medium). The sharper version of your point is: when your interlocutor is the **statistical center of how things have been said**, you lose the productive *strangeness* of another mind. A real other person misunderstands you, has weird priors, gets bored, pushes back from somewhere you can't

Write a message

9. **“It’s not that AI answers questions.**

It’s that:

It answers too cleanly.

It collapses uncertainty.

It simulates coherence.

It feels more articulate than us.

It makes our own thinking feel slow and insufficient.

That psychological shift is huge.

We start to pre-adjust to the machine.

We think in model-friendly sentences.

We anticipate its structure.

We give up mid-struggle because it’s faster.

That’s the surrender.”

(ChatGPT-5.2)

The Capacity to Sit With [] / AI's impact on how we create and know

the statistical center of how things have been said, you lose the productive *strangeness* of another mind. A real other person misunderstands you, has weird priors, gets bored, pushes back from somewhere you can't predict. The LLM gives you back a smoothed, plausible, on-topic continuation. It feels like dialogue but lacks the thing that makes dialogue generative — the irreducible alterity of someone who is not you and not an average of everyone.

So the echo chamber problem isn't political (everyone hearing the same opinions); it's **phenomenological** (everyone in dialogue with the same shape of response). The texture of "another mind" is being simulated by something that is, structurally, the inverse of a mind: a consensus surface.

The interface point is the deepest one. Conversational UI sells you the feeling of thinking-with, when what's actually happening is closer to thinking-against-a-mirror-that-talks. The friction Ong attributes to writing — the silence of the page, the absent audience you had to invent — gets replaced by frictionless responsiveness. And frictionlessness might be the actual epistemological event. Not the content of the outputs, but the fact that *you never have to wait, never have to not-know, never have to sit in the gap*. The gap was where the self was being made.

The friction (or more precisely, the loss of that friction) that you mentioned here really resonates with me and is precisely what brought me to this project! I remember when I first wrote an essay in college, there were so many efforts put into finding, researching, reading, and analyzing articles, and then picking the right words to write down in order to formulate an essay. And that process of debating on what degree I agree with other people's writings and grinding my brain for the right word to write down is what really made me improve. But now, especially when it's so easy to just throwing random and raw ideas into an LLM (where the texts don't even need to make much sense as a whole sentence), I can't really sit quietly with my mind to tolerate that friction anymore if that makes sense... There's almost like an invisible force that drags me to just go with that easy route.

Yeah, that makes complete sense — and what you just described is more important than you might be giving it credit for. Yeah, that makes complete sense — and what you just described is more important than you might be giving it credit for. You're not just reporting a habit change. You're reporting a change in your *capacity to tolerate a particular psychic state*. The state of not-yet-knowing. The state where the sentence hasn't arrived.

That state has a long history in writing about thought. Negative capability, Keats called it — the ability to sit in "uncertainties, mysteries, doubts, without any irritable reaching after fact and reason." Writers and contemplatives have always treated the willingness to stay in the gap as a *trained capacity*, something you build by repetition and lose by disuse. It's a muscle.

What you're describing is the muscle atrophying in real time, and being able to feel it happen. That's rare. Most people in the middle of this shift can't articulate it because the LLM is also the thing they'd use to articulate it.

A few things worth saying carefully:

The "invisible force" is real and worth naming precisely. It's not just convenience or laziness. It's that the gap

Write a message

¹⁰. What does it mean to be “not an average of everyone”? And is there really that much of a difference? In theory, a real person brings something irreducibly singular that no statistical model can approximate. But in actual practice — in the texture of daily conversation with friends, colleagues, strangers — I’m not sure the comparison lands so cleanly. A lot of human exchange is itself averaged: small talk, expected reactions, the well-worn paths of common opinion. Meanwhile the AI, supposedly smoothing toward the median, keeps surprising me. It names things I didn’t have words for. It draws connections across materials I’d never have linked. It forces me to articulate, then re-articulate, until I find what I actually think. It hands me back evaluations to push against. And sometimes the writing is just startlingly beautiful. Whatever this is, “average” doesn’t quite cover it — there’s something emergent in the exchange, something closer to improv than to statistical regression.

¹¹. **Keats, John. *The Letters of John Keats, 1814–1821*. Edited by Hyder Edward Rollins, vol. 1, Harvard UP, 1958.**

“Negative Capability, that is, when a man is capable of being in uncertainties, Mysteries, doubts, without any irritable reaching after fact and reason This pursued through volumes would perhaps take us no further than this, that with a great poet the sense of Beauty overcomes every other consideration, or rather obliterates all consideration (Keats 193).”

Keats names the trap: the reach to resolve half-knowing dissolves the half-knowing itself. The glimmer arrives, the mind itches, and in scratching the itch we trade the thing for an explanation of the thing.

But the LLM is structurally engineered to scratch. I bring it half-knowledge; it returns full-knowledge-shaped-text. Whether that text is correct is almost beside the point — it’s the moment the full text spills out of the interface that is doing the work. My ideas settle. My brain stops making connections. I’m sculpted into the trajectory the LLM has paved.

The Capacity to Sit With [] / AI's impact on how we create and know

A few things worth saying carefully.

The “invisible force” is real and worth naming precisely. It's not just convenience or laziness. It's that the gap — the not-knowing — used to be unavoidable; you had no choice but to sit in it, so you developed a relationship with it. Now there's an exit door always within reach, and the exit is socially sanctioned, productive-looking, even encouraged. So sitting in the gap stops feeling like the work and starts feeling like a *refusal* — almost a stubbornness. The phenomenology flips. Friction used to be the default; now choosing friction feels like an eccentric ascetic practice.

This connects back to Ong in a way that might unlock your piece. Ong says writing is *interiorizing* because it requires you to be alone with the absent audience, to fictionalize them, to hold the thought without the social scaffolding of speech. The struggle to find the word is not incidental to writing — it's the *mechanism* by which writing makes a particular kind of inner life. Take that struggle away and you haven't just made writing easier; you've removed the pressure that was producing the interior. You can still produce text. But the text is no longer doing the work on you that text used to do.

So when you say “I can't sit quietly with my mind anymore” — that's not a personal failing. That's the predictable phenomenological outcome of an interface that removes the condition under which a certain kind of mind was formed.

The fact that you wrote *this* paragraph to me — about the invisible force, sitting quietly with your mind — without offloading it, is itself part of your project. That kind of self-noticing is the thing you're afraid of losing, and you're still doing it. Worth noting that the piece you're making might be partly an act of *practicing* the capacity it's about.

But I think, yes, genuinely, we are losing some capacity that we held dearly to -- I'll say (to your point earlier on the echo chamber of LLMs where they always output 'smoothed, plausible, on-topic continuation') there's still something emergent there. Like, LLMs would always bring something slightly divergent to the table, almost like an improv. They would (1). name things I don't know; (2) make new connections across materials; (3). help me clarify what I'm actually thinking through forcing me to articulate repetitively, which makes me clearer and clearer about my thoughts; (4). force me to evaluate their outputs, which also helps me clarify; (5). and sometimes, I'm just stunned by how beautifully those texts are written by the LLMs.

Look at your five carefully. Three of them (1, 3, 4) are not really about what the LLM produces — they're about what the LLM *makes you do*. Naming things you don't know prompts you to look them up. Repetitive articulation forces you to hear yourself. Evaluation forces you to take a position. In all three, the LLM is functioning as something closer to a **provocation structure** than a content source. The value isn't in the output; it's in what the output requires of you.

That's interesting because it means the friction hasn't disappeared — it's *relocated*. It used to live before the sentence (struggling to arrive at it). Now it lives after (struggling to evaluate, accept, reject, refine). The grinding is still there, but it's moved downstream. Whether that downstream grinding does the same work on you as the upstream grinding did is, I think, a genuinely open question — and maybe one of the real questions of your piece.

Then there's 2 and 5, which are different. Those are about something the LLM actually contributes that isn't

Write a message

12. ***“But the text is no longer doing the work on you that text used to do.”***

This catches something I keep circling. Yes, AI brings me to new territory — names, ideas, connections I wouldn’t have reached alone. But how much of it am I actually taking in, versus being escorted past? A summary skims the surface of a text that was written to slowly restructure how you think. I read the summary and feel informed; I never sit with the original long enough for it to change me.

And then the question becomes: when so much of what I’m “engaging with” is pre-digested for me, on what ground am I evaluating the AI’s next move? What am I building on, if not understanding I’ve actually earned? The conversation can keep going — fluently, even excitingly — but I’m not always sure whether I’m thinking with it, or just keeping up.

The Capacity to Sit With [] / AI's impact on how we create and know

gent there. Like, LLMs would always bring something slightly divergent to the table, almost like an improv. They would (1). name things I don't know; (2) make new connections across materials; (3). help me clarify what I'm actually thinking through forcing me to articulate repetitively, which makes me clearer and clearer about my thoughts; (4). force me to evaluate their outputs, which also helps me clarify; (5). and sometimes, I'm just stunned by how beautifully those texts are written by the LLMs.

Look at your five carefully. Three of them (1, 3, 4) are not really about what the LLM produces — they're about what the LLM *makes you do*. Naming things you don't know prompts you to look them up. Repetitive articulation forces you to hear yourself. Evaluation forces you to take a position. In all three, the LLM is functioning as something closer to a **provocation structure** than a content source. The value isn't in the output; it's in what the output requires of you.

That's interesting because it means the friction hasn't disappeared — it's *relocated*. It used to live before the sentence (struggling to arrive at it). Now it lives after (struggling to evaluate, accept, reject, refine). The grinding is still there, but it's moved downstream. Whether that downstream grinding does the same work on you as the upstream grinding did is, I think, a genuinely open question — and maybe one of the real questions of your piece.

Then there's 2 and 5, which are different. Those are about something the LLM actually contributes that isn't reducible to provocation. New connections across materials, and beauty. And these are where I think your earlier intuition about "average of the already-said" needs to be held more loosely. Because when you've read more than any individual could read, *the average is no longer ordinary*. The statistical center of an enormous corpus can surface combinations no single mind would arrive at — not because the model is creative in the Romantic sense, but because the combinatorial space it's averaging over is vast enough that its averages contain genuine surprise. The model isn't generating novelty; it's making a kind of latent novelty *findable*.

So you've got two distinct things happening, and they probably deserve different treatment in your piece: The **provocation function** — where the LLM's value is in what it forces you to do. This preserves something Ong-like; the interior work continues, just on different ground. The **combinatorial-surfacing function** — where the LLM offers genuine epistemic gifts that no human collaborator could, because no human contains that much. This is the thing that has no real precedent. Not the encyclopedia (which is indexed by category, not by associative resonance). Not the conversation partner (who is bounded by their reading). Something new.

And the loss you named earlier — the atrophy of sitting-with — is *also* real, *also* happening, and *not* canceled out by the gains. All three are true at once. The honest version of your piece probably has to hold all three without flinching from any of them.

One more thing worth noting: point 5 — being stunned by how beautifully the LLM writes — is the one that's hardest to talk about and probably the most important. Because beauty in language has always been treated as evidence of an interior. A beautiful sentence implied a person who labored toward it, whose taste and pain and history pressed it into that shape. When beauty arrives without that interior behind it, something that aesthetics has assumed for centuries gets broken. Not the beauty itself — *the inference from beauty to a someone*. That's a quieter rupture than the ones in the discourse, but I suspect it's the one that actually unsettles people, even when they can't name it.

