Think More, Not Less

Design for Human Creativity with AI

Glenn Fajardo, Fall 2026
Stanford Accelerator for Learning
AI in Teaching and Learning at Stanford: Thought Leadership

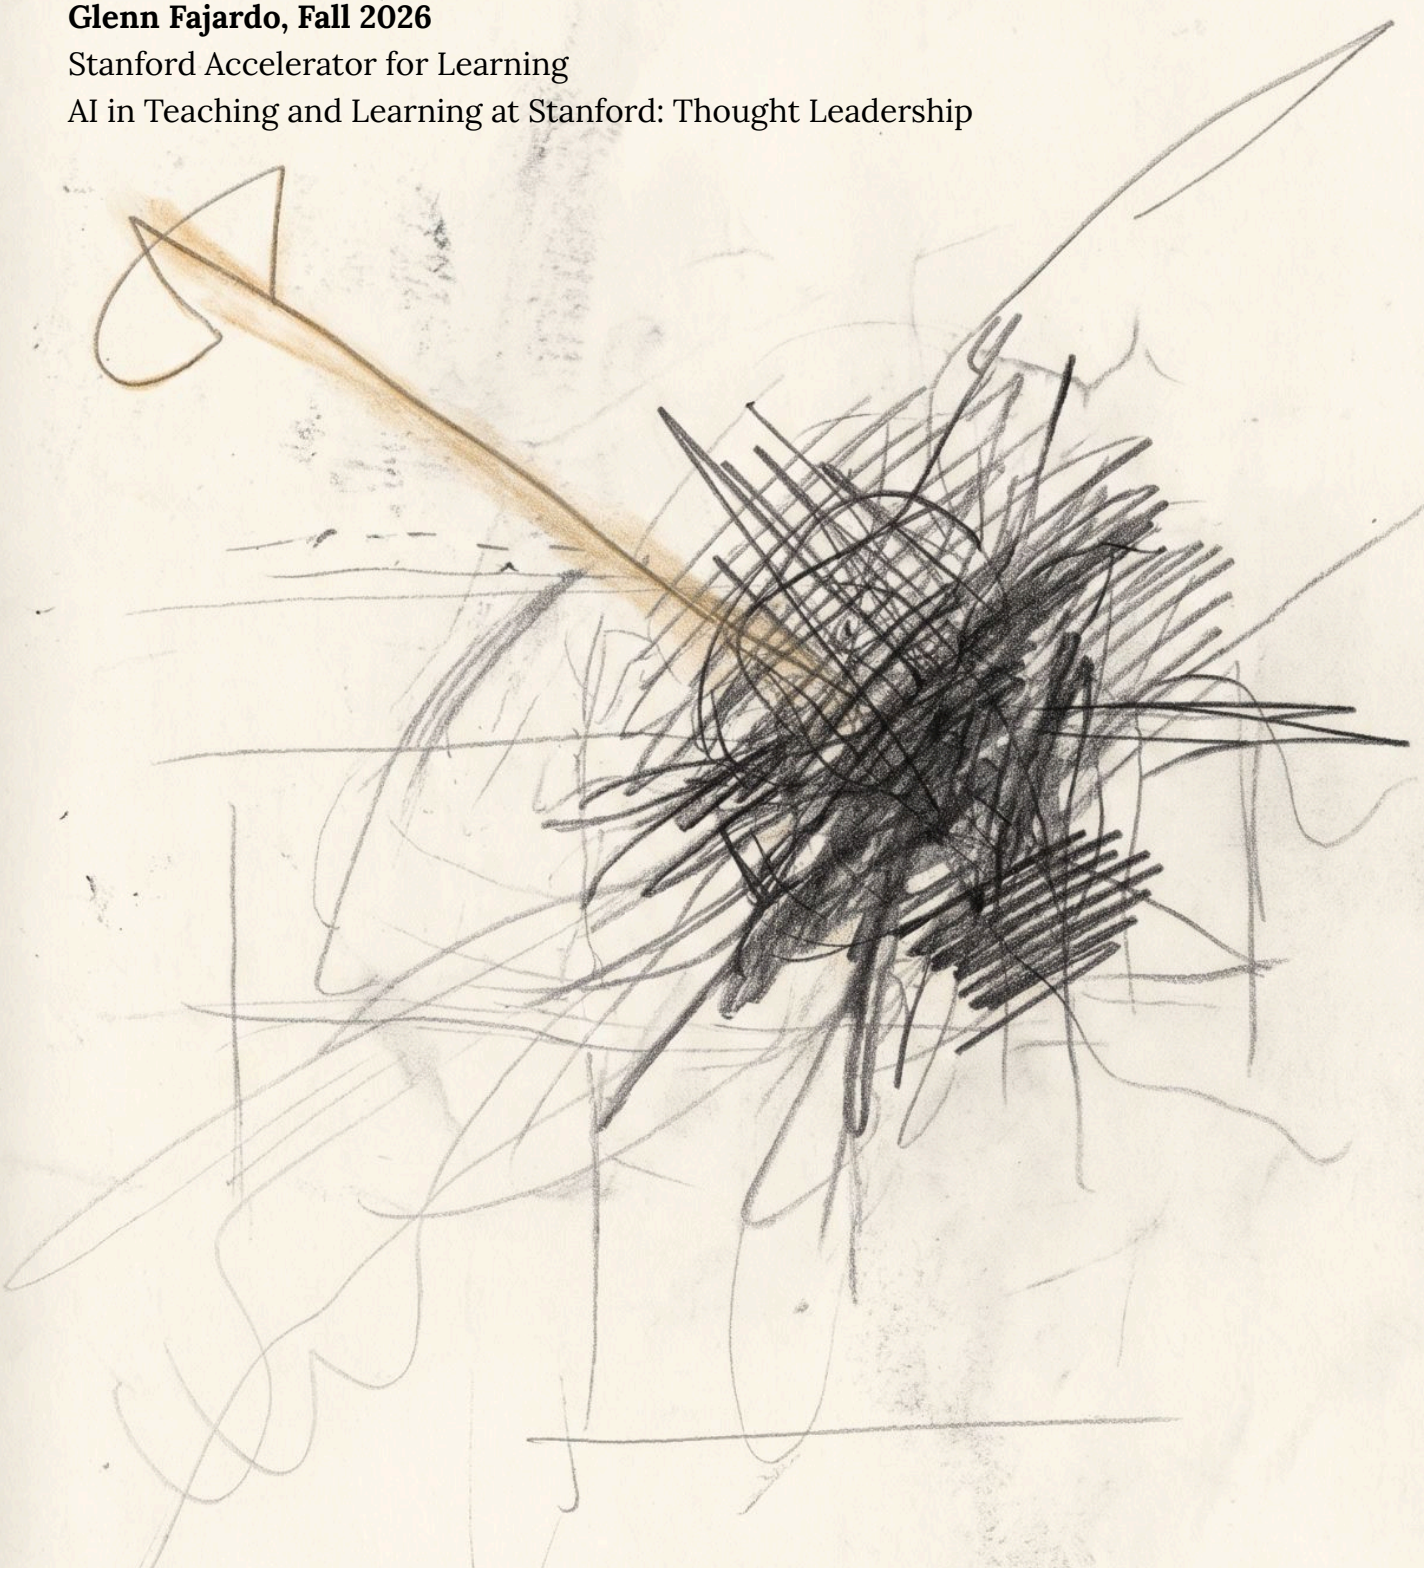

Contents

Introduction: Forgetting our curiosity? 2

1. The Gravitational Pull of the Most Likely 6

2. The Locus of Choices Changes 12
   Something to try: Locus of choices 18

3. Solving the Wrong Problem, Faster 19
   Something to try: Four frames 24

4. Bring Your Intent 25
   Something to try: Write down what you're after, then predict 30

5. Playing with Constraints 32
   Something to try: Different directions with constraints 38

6. Noticing Yourself and Your Materials 40
   Something to try: Name your reaction 46

7. A Mood Board, Not a Menu 47
   Something to try: A mood board, not a menu 55

8. Who Can Work This Way? 56
   Something to try: Go upstream once 61

9. Not a Solo Sport 63
   Something to try: A writers' room with AI in it 69

Conclusion: Remembering to Be a Designer 70

Acknowledgments 71
Bibliography 72

Introduction: Forgetting our curiosity?

I've been a long-time member of the Stanford d.school's teaching community, and in Spring Quarter 2026, I was supporting the d.school Home Team with AI experimentation. Each person was conducting their own personal experiment with AI, to play around with it and see what they could do.

A very, very skilled designer was the first to raise their hand, eager to get coaching, and I was happy to help and share in the curiosity. They were running an experiment using an AI agent to detect signals around how people discuss AI's impact on the design field. (Yes, a little lower case "meta.")

I thought it was a nice fun-size experiment, clearly doable but with some nuances that could be interesting to work through. After listening to their overview of what they were trying to do, I asked what I thought were pretty basic design questions, for example: "What different ways have you thought of to frame the problem you're trying to solve?"

After the person heard my questions, they paused, stared for a second, and then the room suddenly got brighter as if a light bulb lit up over their head. "OMG, I TOTALLY FORGOT TO BE A DESIGNER!" (Ok, maybe the brightness was a shift in sun through the window.)

This very, very skilled designer reflected on how they had jumped to go "fast" and "efficient" because “that's what people do with AI.” And it's true. That is how a lot of people do AI. That’s the dominant mental model. But it doesn't always have to be that way.

I’ve kept thinking back to my colleague who forgot to be a designer. The good news is that kind of forgetting can be momentary. It took one question to bring a very skilled designer back.

…

In my mind, I often put air quotes around the word "productive." As in, "Generative AI can make us more 'productive.'" It's not that I doubt the potential utility of AI. Not at all!

To me, it's roughly like saying a smartphone is "useful." A smartphone is obviously useful, but there's some stuff (or "stuff") that comes with that usefulness. Or it's like saying a TV is "entertaining." There's stuff that comes with that.

I'm curious about how we can use technologies more thoughtfully and intentionally, but I know it's also easy to default to defaults.

And I've wondered **why.** Why is it so easy to default to defaults? And I've wondered **how?** How can we bring out our human creativity when we're working with AI in the mix?

I also, quite frankly, find myself glazing over when I hear the words “taste,” “judgment,” and “discernment.” For example, "As execution gets easier, judgment matters more." I’m not disagreeing with the premise, but it raises the question, *how* does one become better at judgment? I don’t think it’s helpful for everyone saying you need “taste” when no one is talking about *how* to develop your own.

So I've been exploring these questions as a design educator. I know it’s awfully convenient for a design educator to say this, but the more I explore these questions around human creativity and AI, the more convinced I am that the mindsets and skills of design are becoming more essential for creativity in the age of AI. And I think those mindsets and skills might need some **radical evolution**.

“Radical evolution? WTF is that? Isn’t that a contradiction?” Not necessarily. I’m thinking of “radical” in the sense of going back to the foundation. And I’m thinking of “evolution” in the sense of a gradual process of change that doesn’t suddenly throw away the fundamentals.

The idea of radical evolution is what happens when we resist the pull toward simple stories and “THIS CHANGES EVERYTHING,” and dial up curiosity instead.

When we’re more curious, we realize that academic research is more interesting than we think, and largely because of the nuances.

One of the people I interviewed, Harsh Kumar, co-wrote a paper that went viral, “Human Creativity in the Age of LLMs: Randomized Experiments on Divergent and Convergent Thinking.” I myself remember being told by one of my students about the social media discourse: “There’s a paper that shows that AI coaching doesn’t work!”

I think Kumar was both bemused and a bit frustrated:

> *"What happened was, we didn't design it to be fully powered enough to make very strong scientific claims. That's why we submitted it to CHI: to start this dialogue about conducting more experiments longitudinally, with this being the first experiment… I think as it went to social media and people interpreted the findings, they focused a lot on just the negative framing of the paper, which we had mentioned as one potential side of things, where they can go wrong. But there's also the other side, that maybe it can help us be better… I started getting so many messages on Instagram and TikTok and everywhere, people asking, 'How should I use this?' And I was like, 'we don't know!' That's the whole point of the paper: we should do more research to find more scientifically rigorous ways to find out what makes a good coach that we can design for interaction."*

Wanting a simple answer is understandable. And there are times when AI-as-easy-button actually is the sensible thing to do!

My friend told me a story about having to cancel a patent, and how she thought it was so awesome to use AI to do most of the process, so that she didn't have to think about it. She doesn't need to have an ongoing understanding of the patent cancellation process, nor does she need to be engaged in a Socratic dialogue about it. She just needed to get it done.

And I'll admit, sometimes I want the easy button when it *isn't* the right call, like when I'm writing a tricky email, the kind where I first have to figure out what I actually think, and then how to say it. I'm getting better at noticing that feeling, though I fail to notice sometimes.

I wonder if how we often talk about "cognitive offloading" in the context of AI ends up misframing the problem, as in, "Oh no, cognitive offloading!" Writing down your grocery list, even if it's on a piece of paper, is a form of cognitive offloading.

I think where we can get into trouble with cognitive offloading and AI is when we use it with no real understanding of what we are doing, and when we aren't putting any of our own ideas, perspectives, or personal input into the work. I think we need to be more curious about *that*. (For example, the false sense of mastery that comes from generating an argument from a rough outline without actually reasoning through it.)

We need to be more curious, not less, if we want to think more, not less.

I became a curiosity sicko for the last few months and did a lit review of 145 sources and interviewed 24 people. (Because curiosity. That's the tweet.) The sources were mainly academic papers, including preprints, many of which were part of the flood of AI papers that no one can keep up with, but also many of which were not: research on automation from the last four decades, research on the nature of creativity, research on the impacts of the use of non-AI technologies from Google search to the invention of the camera. The 24 interviews were with designers, artists, writers, educators, and researchers from all over the world. I asked them about not only their perspectives from their current work, but also what they've drawn from their past lives as ballerinas, b-boys, and colorists.

More thinking sounds like work, and it is! Why do it in the age of AI? Four reasons are top of mind for me.

- First, AI is making it easy to solve the wrong problems quickly. It's easier to feel like you're done when in fact you haven't done anything meaningful.
- Second, better thinking can lead to solutions that are usefully different from what everyone else is suggesting.
- Third, you can feel a greater sense of personal agency and ownership.
- And fourth, you can get closer to the desirable psychological state of flow, which Mihaly Csikszentmihalyi defines as the state where the level of challenge you take on meets the level of your skill.

So let's get curious, starting with this question I mentioned: Why is it easy to default to defaults?

1. The Gravitational Pull of the Most Likely

Let's say we were trying to come up with an idea for a business… an *innovative* business. We'll prompt ChatGPT with this:

> *You are Steve Jobs. Create 5 innovative business ideas.*

What do we get? When I tried this in ChatGPT with a chat that was not connected to my personal ChatGPT memory, I got these ideas:

- "Life OS" ("One AI that manages your entire life")
- "Second Brain" ("Your life becomes searchable")
- "Zero Wait" ("Shopping becomes almost invisible")
- "Micro University" ("A personalized education")
- "Digital Twin" (an AI employee for small businesses)

Not terribly interesting. But what happens when I try this in ChatGPT in a chat connected to my personal ChatGPT memory? Then I get:

- "The Second Brain That Makes You Smarter"
- "The School of Things That Don't Exist Yet"
- "The Anti-Algorithm"
- "The Personal Manufacturing Company"
- "The Human Potential Studio"

Still not very good or particularly original ideas. Many of you reading this won't be surprised by that. But if we dig just a bit more, we find a few interesting nuggets.

For example, a second brain appears in both runs. Here, personalization didn't change the category; it just changed the framing. The generic second brain makes "your life searchable"; the connected-to-memory one "helps you discover what you think." It's the same structure painted over with different values. The generic one prizes convenience, efficiency, and invisibility. The one connected to my personal memory prizes curiosity and expanding what people can do. Same basic idea though.

(Sidenote: Fun to get an idea for an "anti-algorithm" idea from an algorithm!)

This is not a scientific experiment here, but it does start our journey to the gravitational pull of the most likely. What happens when you do something like this over and over?

Columbia Business School researchers Yuting Deng, Melanie Brucks, and Olivier Toubia explored a paradox in generative AI: while individual AI-generated ideas often appear novel, outputs across separate sessions rapidly converge, lacking the collective diversity produced by human groups.

Part of the pull to the most likely has to do with the fundamental behavior of large language models, the behavior of most likely prediction. Models are built to produce what's most likely given what came before, whether that's the next word or the next detail in an image.

Some people have the intuition that what you need to do is increase the temperature of the AI. Temperature, in the context of generative AI, controls how much "creative risk" the AI takes when choosing what to generate. But it turns out this isn't a great approach. As writer and co-founder of Sudowrite James Yu notes, "temperature helps a little, but it doesn't solve it." The challenge is that as you turn up temperature, things can go so far afield to the point of not being helpful.

Any single AI session can produce ideas that seem fresh on their own. Run many sessions, though, and the ideas keep landing in the same places.

As Deng, Brucks, and Toubia note, models "aggregate knowledge into a unified distribution rather than exhibiting the knowledge partitioning inherent to human populations, where each person occupies a distinct region of the knowledge space." A crowd of people has different perspectives. A model blends them into one average. In some ways, when mixing things, a model behaves less like a salad or even a melting pot stew. It might be closer to a smoothie. The result is called **homogenization**: ideas and outputs become more alike across people.

The same thing can happen to the people using AI. Researchers Anil Doshi of University College London and Oliver Hauser of the University of Exeter provided writers with AI-generated ideas for short stories. Independent evaluators judged the resulting stories to be more creative, better written, and more enjoyable, with the greatest quality gains seen among writers who had lower baseline creativity. However, the stories also became significantly more similar to one another, leading

the researchers to warn of “an increase in individual creativity at the risk of losing collective diversity.”

It shows up at the scale of a whole field, too. Analyzing 41 million scientific papers, mostly from before generative AI, researchers from Tsinghua University and the University of Chicago found that scientists who do use AI tools publish more, get cited more, and lead projects sooner. But across science, AI adoption narrowed the range of topics studied and reduced how much scientists engaged with one another’s work: “an expansion of individual scientists’ impact but a contraction in collective science’s reach.”

James Yu has seen this firsthand:

> *"go to almost any model and say, please write me an opening for a small-town romance. There's a ninety percent chance the town will be called Willow Creek. We don't know why... It shows up everywhere: Claude, OpenAI, DeepSeek, all of them. Somehow, in the branches of training, Willow Creek has been ordained."*

And as AI director and artist László Gaál notes, “In a film, you don't want everyone to look like a Korean skincare product commercial,” which the commercial tendencies of image and video generation tend to drift towards.

But wait… there’s even more feeding the gravitational pull! In addition to model behavior, there’s another kind of behavior that powers the pull to the most likely. Human behavior.

…

Several aspects of human behavior contribute to that gravitational pull. One is that alongside the most likely prediction of AI models, there's most likely prompting by humans (or at least *more* likely prompting).

For example, many people who try to use AI to help them come up with a business idea ask the AI to behave like Steve Jobs or another creative entrepreneur. Does that approach lead you to different ideas?

Deng, Brucks, and Toubia explored this aspect with different prompting approaches, with a goal to come up with "ideas for a fitness product that would help people improve their fitness." It turns out you get significantly higher collective

idea diversity for a new business if you prompt large language models with varied, "ordinary" personas (such as a dairy farmer, an ink chemist, a museum curator, or a college student) than you do by prompting them with "creative entrepreneur" personas (such as Steve Jobs).

What this suggests is not that “act like a dairy farmer” is a singularly secret prompt. (Tech bro voice: “Duuude… just tell it to act like a dairy farmer!”) Instead, what it suggests is that if your goal is to come up with a distinctive idea, you may be better off avoiding where others are likely to look, and instead exploring a wide variety of possibilities.

Another aspect of human behavior that can contribute to the gravitational pull of the most likely might be called **“the quick good enough.”** A16z design engineering fellow (and former ballerina) Rina Takikawa describes this pattern:

> “*The first pass looks pretty decent; let me move on to the next. And then it goes into a loop, an autopilot. That's how people's default starts changing, because it becomes mental muscle memory.*”

Good enough can reduce our feelings of uncertainty, and **our desire to reduce uncertainty** is another facet of human behavior that adds to the gravitational pull.

Organizational behavior researchers Jennifer Mueller, Shimul Melwani, and Jack Goncalo shared their findings on this in a 2012 paper called “The Bias Against Creativity: Why People Desire but Reject Creative Ideas.” People say they want creative ideas, and they mean it. But when they're feeling uncertain, they carry a hidden bias against them. The more someone wants to reduce uncertainty, the harder it gets for them to recognize a creative idea when they see one. They rate highly innovative ideas as less creative and less desirable.

And there’s **fixation**. In a seminal 1991 paper, design researcher David Jansson and cognitive psychologist Steven Smith described design fixation as a blind adherence to a set of ideas or concepts limiting the output of conceptual design. People unknowingly copy features of an example or a past solution, even features that don't fit the problem in front of them, and it keeps them from exploring other, possibly better, directions.

Fixation can get amplified by how people go about using generative AI. University of Melbourne human-computer interaction researchers Samangi Wadinambiarachchi and her colleagues studied how AI image generation (Midjourney) can affect what designers do in a visual ideation task.

If I were writing an attention-grabbing headline for their study, it'd be something like: "AI tools intended to make people more creative can make them less creative." And that'd be a mostly accurate headline. As the authors state, "Participants who used AI produced fewer ideas, with less variety and lower originality compared to a baseline."

But I think what's more interesting is *why*.

First, designers frequently relied on prompts containing keywords copied directly from the design brief. Before we get on the study participants for being lazy, let's take a step back and think about it: Using instructions from a design brief isn't a crazy thing to do! That kind of behavior might be analogous to how you use AI at times. And more strikingly, the designers in this study were *unaware* of how their own prompting habits were anchoring them to the initial example provided in the design brief. During interviews, participants explicitly claimed they were trying to avoid copying when creating their own work.

Second, when designers used AI to generate images as inspirations for their own drawings, participants ended up fixated on the images generated by the AI and copied what they saw. For example, one participant talked about the tweak-a-bit approach that many of us use when using AI:

> "*It helped a lot of the inspiration for a lot of the designs that very much I just put down what I wanted it to give me, and I would just kind of copy it and then tweak it a little bit for the designs.*"

This approach led to less originality and variety in the designs. Tweak-a-bit has its uses, but when everyone is doing it a lot, everything starts to look the same.

This isn't only an AI story. Justin Keoninh, a designer who has been a b-boy for decades, remembers when you could watch a dancer and say, "that person's from Miami, or that person's from New York." Then YouTube came along, and "everything started looking very similar… you couldn't tell where anybody was from."

When I asked TU Delft design researcher Sonja Rozental why even creative colleagues seem less creative with AI, she wondered:

> *"Maybe it's because there's this underlying assumption of productivity tied to AI-driven technologies, that we use them to be more productive, meaning we get the task done quicker. And when we start with that mindset, I think it's very hard for ideas to emerge."*

One research scientist at an AI lab put it this way: "Productivity is really about how do you get to an answer as quickly as possible."

And as James Yu put it, "you steer the AI, but the AI also steers you."

Speed has its strengths, and productivity has its place. But in order for us to do creative work with AI in the mix, we need to recognize the gravitational pull of the most likely. And then we can start to see how we can make different choices.

2. The Locus of Choices Changes

2025 was supposed to be the year of AI agents, but it felt like it was this year, 2026, where I would hear more and more about people managing multiple agents in their work throughout the day. (Perhaps you are one of these people!)

I've been very curious about what it feels like to work this way, when it feels like your job depends on it. So I asked people about their experiences.

I asked Rina Takikawa if she spent part of her day working with multiple agents simultaneously, and she said, "No, it's like *every minute* I'm talking to ten agents! I'm working on ten different things at once." What does that feel like to her? "Lots of mental switching."

Earlier in our conversation, she described the pressure behind it:

> *"It's very difficult now to think and operate like a designer. Because as a designer, you need space in your brain to really tinker... when you're in this environment of people shipping multiple PRs a day... you feel the need to catch up to that. And that becomes almost the new baseline."*

I also asked Gloria Chua, a design manager at Singapore's Ministry of Education, what working with AI feels like. Describing how she makes illustrations, she said:

> *"I would prompt it and hope that something works. And if it doesn't, then you just try again. It feels like a spinny jackpot kind of thing. You just don't know what you're going to get."*

Later in our conversation, I described what I'd been hearing from people managing many agents at once: it's like playing a bunch of slow-motion slot machines. You don't totally know what you're going to get, and waiting becomes a big part of it. It feels really weird. But when you "win," you do get this hit, this feeling of productivity. Chua recognized it right away:

> *"I love the analogy of multiple slot machines. One of the aspects it made me think of is that I spend a lot of time waiting. And it sucks, because then I have to multitask, and I really hate jumping between things. It's really, really frustrating."*

I replied to Chua:

> *“I've been experimenting with whether there are things I can do while I'm waiting, other than staring at the screen going, "Are you done yet? Are you done yet?”*

The psychologist Sophie Leroy has a name for what makes this hard: “attention residue.” When we switch away from an unfinished task, part of our attention stays behind, and our work on the next task is impacted.

…

If you look over the history of technology, there's a familiar story: new technology X handles the routine parts, so humans can focus on the high-value work. But it's rarely that simple.

A new technology does not simply replace a human weakness. It creates new strengths and new weaknesses, and these are often hard to anticipate.

Lisanne Bainbridge’s 1983 paper, "Ironies of Automation," is a classic in automation research. In it, she describes several paradoxes inherent in trying to design systems with more automation.

One is the leftover paradox: Automation easily handles predictable, standardized routines, leaving the operator to handle the edge cases that designers could not figure out how to automate. The human is left with an arbitrarily fragmented set of the hardest, least predictable problems. People’s jobs with the automated system can actually end up feeling worse.

Starting in the late 1980s, David Woods and his collaborators repeatedly argued against the "technology-as-replacement" paradigm, demonstrating that automation redistributes workload clumsily rather than simply subtracting human labor.

I think it’s worth keeping these patterns in mind today. In a 2024 paper that was a nod to Bainbridge, “Ironies of Generative AI: Understanding and Mitigating Productivity Loss in Human-AI Interaction,” technology and human-AI interaction researchers Auste Simkute and her colleagues identify four sources of loss that can come into play: a shift in the user's role from production to evaluation, unhelpful restructuring of workflows, interruptions, and a tendency for automation to make

easy tasks easier and hard tasks harder. They also note, “Humans are also ineffective in shifting cognitive resources saved by automation to support more difficult tasks.”

And there’s also another pattern in history. **When a new technology emerges, what often happens is the locus of human choices changes. How people do things changes, and so does where they make choices.**

Often, it’s not obvious, and a new technology can appear to be just dumbing people down.

In his 2006 article "Myths and misconceptions: photography and painting in the nineteenth century," art historian Hans Rooseboom challenges the widespread assumption that the invention of photography immediately devastated traditional painting. Rooseboom recalls the story of Tholen’s devil in 1902:

> *The Dutch painter W.B. Tholen wrote to an art critic that "photography is steering painting in such a wrong direction that it's time this was stated out loud." He pictured the struggling artist visited by the devil: "but why go to all that trouble, why tire yourself out with studying and drawing? Look here. And from his pocket he takes a photo of the scene that one is trying to depict. It contains everything you need, ready-made and absolutely perfect... People will be amazed, it will make you rich and respected, you can make what you like without any special knowledge. I will tell you the secret, says the devil, on condition that your heart belongs to me."*

Tholen did have a point: if you were to skip the studies of painting and drawing, you would have a different relationship with the creation of an image. But Rooseboom found photography wasn’t stopping people from painting. Many well-known Dutch painters of the 1800s, including Breitner, Israëls, Mesdag, and Schwartze, regularly used photographs to help them study their subjects. They didn’t usually advertise this, since art critics might have looked down on it. (Hide your AI use, anyone?)

Critics said photography couldn't be art because "the role of technology, physics and chemistry in making a photograph was too great, the role of intellect and hand too small."

But photography itself eventually became a creative medium too. Over time, photographers like Ansel Adams didn't just click a shutter. They made deliberate, expressive choices with aperture, contrast, and framing. The locus of creative choices changed relative to painting.

The same happens with music. In 1973, at his sister's back-to-school party in the Bronx, DJ Kool Herc used two turntables to isolate and repeat the break in a funk record. As the writer Stephen Marche puts it, Herc "was in a sense doing the most derivative thing possible: He was replaying the good bit from a hit record over and over again. He was also doing the most original thing possible: He was inventing the break." It took years for many people to see the choices and skill involved: hearing which pieces of sound could be transformed and recontextualized into something new.

I think we're still very much in the process of figuring out the nuances of how we make choices with AI in the mix.

As James Yu told me, with AI "there's a new entity, a non-human, coming into conversation using language, for the first time in a fluent manner, and in a manner you can collaborate with. We really haven't had that kind of non-human language collaborator, ever."

So with this new conversational entity involved, where might the human choices be going? In my interviews, they kept showing up in a few places.

One place is in **curation**. Punya Mishra, a professor and educational designer at Arizona State University, told me:

> "*If you're an expert in a domain, the kind of work you do has shifted. You've gone from being a person who creates the first draft and edits it, to somebody who curates multiple drafts to get the best possible draft. That's a different kind of role... It's become more of a managerial role than a creative role… and it's got its own kind of load. People talk about efficiency, but I'm not sure it works that way.*"

He'd felt this firsthand, working with Claude on a speech for his son's wedding:

> "*At the end of the day, every word, everything, had to be okayed by me. Even if there was a phrase Claude put in, and I said, that's actually good, I like it: I have*

> *to like it. It has to be something I would say. And cultivating that is a lot of work... It's a different kind of effort than when I sat and struggled over writing a paragraph. Now I have three versions of that paragraph, none of which seem right.*"

The novelist Stephen Marche, who experimented by producing a novella with ChatGPT and Sudowrite, sees curating as a skill in its own right: "You need more understanding of literary style, not less... The creator of meaningful literary AI art will be, in effect, a literary curator." (His essay on the experiment is titled "The Future of Writing Is a Lot Like Hip-Hop.")

A second place where human choice shows up is in **designing process**. Sonja Rozental told me artists like Jon McCormack "become meta-artists, in the sense that instead of drawing themselves, they design the very process." McCormack himself told me his concept of "process aesthetics":

> *Maybe you've heard mathematicians talk about maths as having some kind of beauty, like an elegant solution. It's related to that, I think. The aesthetics of a process is not visual; it's in the process itself. How beautiful is the process in achieving its outcome? It's similar in computing as well: people talk about elegant algorithms that, in a small amount of code, do something really profound, in a mathematical sense. It's strongly related to that.*

And a third place where choices can reside is in **finishing**. Pann Lim, co-founder and creative director of Kinetic Singapore, shared this perspective:

> *"As of now, and I believe this might change: if you start generating images, you'll know there will be some artifacts, or some inaccuracies, from the prompt you're doing. If you look at the overall image, it's already 85 to 90% good. That last 10% is hard to perfect. You can prompt again and again, and it'll just take you further and further away, because it doesn't know exactly what you're talking about, to finish the last 10%, or the last 5%, which is actually the important part. So in the last mile, the human needs to take over that file. Because to keep prompting to get the last 10% done is even more tedious than "just give me the file and I'll complete it." It'll be done in five minutes, instead of spending an hour and burning credits.*"

Curation, designing process, and finishing are three examples of where human choices might reside, and we might find other places. We don't have to settle this once and for all. We're early in figuring out what this medium asks of us.

It takes time for people to play, experiment, and figure out the new human choices that are possible with a new medium. New mediums mimic old mediums. That continues until people figure out the "grammar" and possibilities of the new.

For example, the close-up shot is one of the most fundamental elements in the grammar of movies. It is one of the main reasons that movies can create emotional impact in the way that they do, focusing audience attention on key details and character reactions. Movies were invented in the 1890s. The close-up shot didn't become widely used until the 1910s. With AI, we may still be before the close-up.

If the choices are moving with AI in the mix, a practical question follows: which choices do you want to keep making, and which new ones do you want to take on?

Remember Simkute's point that people don't automatically shift the effort automation saves toward harder work? Freed-up attention doesn't go anywhere on its own.

In a series of experiments, cognitive psychologist Sandra Grinschgl and a team of researchers from the University of Tübingen had people copy patterns between two windows they couldn't see at the same time. Checking back at the original instead of holding it in memory helped people finish faster, but they remembered the patterns less well afterward. The memory cost was smaller when people knew a memory test was coming. The researchers concluded that the goal matters: offloading "did have detrimental effects on memory without such a goal."

Having an intentional goal in mind can help guide you to what choices you want to make. For example, if you had a goal of "be able to defend this proposal under heavy scrutiny" that might lead you to a different set of choices than "share information about an optional social activity that I'm not going to myself."

When considering where you want to make choices, I think it's also good to remind ourselves that the answer is not to make everything harder!

Vera van der Burg, a researcher who trains her own AI models, has thought a lot about slowness in creative work. "If you just slow stuff down, it's just annoying," she

told me. “That’s not the point... it’s about what you do with it.” The goal isn't more struggle. It's productive struggle in the places that matter to you, where you want to understand and shape the work yourself.

There’s a wrinkle here, though: we don’t always know where we’re making choices. The writer Ted Chiang has argued that “when you are writing fiction, you are, consciously or unconsciously, making a choice about almost every word you type.” James Yu, who has talked with Chiang about writing and AI, characterizes what you’re doing differently. “I think language is kind of a shared unconscious decision mechanism, and also culture, and also storytelling,” he said. “We’re just going down this river when I’m writing these words.” He sometimes uses AI to notice when that’s happening. He’ll ask a model for the most typical way his text should continue, and then do the opposite: “it’s almost using it as a barometer, and then I can decide.” The model’s pull toward the most likely, the thing we’ve been wary of, can also show you where you’ve been drifting along with it.

So the question is where we want to be making choices now. Some of those choices are familiar ones that still matter. Some are new-ish, like the curating, process design, and finishing we just looked at. And some sit at the very start of the work, before we ever open a chat window: what problem we’re actually trying to solve, and what we’re trying to say. That’s where we’ll go next.

Something to try: Locus of choices

Here’s a short activity you can do to start to notice your choices and non-choices when you use AI.

1. First, jot down some notes for yourself on what you think the choices are when you: a) write an email yourself, b) use AI to help you write an email.
2. Then write a prompt for an AI to write an email you need to write.
3. Then write that email yourself.
4. Then generate with AI from your prompt in #2, and then edit the result.
5. Here’s the most important part: Reflect on how the two felt, what choices you made in each, and how the work felt different. It's not about which is better. (There may not be a clear answer to that.) It's about starting to notice where your choices moved.

3. Solving the Wrong Problem, Faster

The Sydney Opera House had a problem that kept coming back. Over the years, protesters had repeatedly climbed its famous white "sails" to unroll banners, and once to paint a slogan on the building, and every climb drew a lot of media attention. In 2005, another incident made it clear that the usual security response wasn't working.

The obvious answer was "even more security." One eminent criminologist suggested fencing off the whole area and charging an entry fee. But that ran into a paradox. The Opera House attracts protesters because it's iconic. And because it's iconic, a UNESCO World Heritage site, it can't be altered. Closing off the podium would mean, in University of Technology Sydney design researcher Kees Dorst's words, that "the Sydney Opera House cannot be fully experienced as a special place."

So the Opera House asked Dorst's Designing Out Crime center for a new way to stop the climbers. Instead of starting with security, the team stepped back and looked at the wider situation. About 8 million people visit every year, mainly tourists from around the world. The headland is also sacred to the Gadigal people of the Eora Nation. Yet Sydneysiders themselves tend to avoid it, and their contact with the building has shrunk to concertgoers "hurrying to the box office." And the building had become a kind of photo cliché, captured in "the millions of nearly identical pictures that are taken of the side view of the building."

Dorst and his team explored alternative ways to look at the situation and frame the problem:

- What if it is a problem of providing liveliness and rejuvenation? Then it might be about enlivening a public space outside the city center, and make it a city square that Sydney residents had a reason to go to on a regular basis.
- What if it is a problem of providing contemplation or spiritual experience? Then it might be about enhancing the natural lighting to create spaces where the audience experiences a special effect.
- What if it is a problem of creating a sense of place, giving meaning to having arrived and being there? Then it might be about providing more depth to what the visitor experiences, through storytelling and opportunities for reflection.

With the new frames, the team imagined the podium alive around the clock, from morning yoga to late-night light art, with floor sensors that could run events and also notice someone lingering at a climbing spot. Constant activity "would make it much harder for people to climb up without attracting attention." These were proposals from a short project, and carrying them out "would require a major transformation." But these multiple frames opened up options a security brief never would have.

Now imagine they had asked AI, 'How do we stop protesters from climbing the Sydney Opera House's sails?' That question is already firmly rooted in the security frame, and would lead you down particular paths.

That question would have arrived as what psychologists Jacob Getzels and Mihaly Csikszentmihalyi call a *presented problem*, where everything is given to you: the problem, the method, and the kind of solution you're looking for are already known. All you have to do is execute.

In a classic 1976 study, they watched 31 advanced art students compose and draw a still life. They contrasted *presented problems* with *discovered problems*, where you don't yet know what the problem is, how to approach it, or what a good solution would look like.

Students who did more exploring while they were still forming the problem produced work that was rated more original, though not more skillfully crafted. And **for the most original artists, the problem often kept shifting as they worked.** As Getzels and Csikszentmihalyi put it, "It would seem that creative work is the outcome of a more or less pure form of the discovered problem-solving process."

The study is correlational, and it was done long before generative AI. But it's a useful lens. When we use AI mainly for faster solutions, we tend to treat creative work as a presented problem: write the prompt, let the model generate, take the polished draft.

László Gaál, who spent about twenty years as a colorist before becoming an AI director, sees this playing out in film. As AI makes short films "very, very similar to writing," he told me, "the main thing will be: do you have something inside your head that would be interesting for other people to watch?" Your film has to be

worth those minutes of someone's life compared with "the ten thousand others that came out that week." His diagnosis:

> "*I don't think we're lacking a better model. We're lacking better thoughts, and better shots, and better lines, and better scripts.*"

So what does reframing actually look like on a practical level?

Justin Keoninh, a product designer who mentors younger designers, used to give his consulting clients an example. "Let's say you're Netflix, and you hire me to design a whole solution for improving the logistics of mailing DVDs." He could design a whole system for it. "But the real question was: how do you get content to your users faster? And what they did was, you know what, we could just stream this. That was the real answer to the question."

And reframing doesn't have to be that big. Keoninh also teaches a startup class where each student gets a hundred dollars and has to make a hundred dollars back from paying customers. Most students start by building an app. "And I'm like, why? What problem are you solving with an app?"

One student he mentored, Charlene, saw that sourcing hair weaves was a real problem in her community. She built a business that, by Keoninh's account, eventually made six figures a month.

> "*And she never created an app. She literally got weaves and sold them out of the back of her trunk. Her biggest problem was distribution and inventory... she understood the problem of her own community and her customers.*"

So… why is it so easy to skip this step with AI?

Abel Valdez, a designer in Mexico City who works with governments on digital public services, described what happens when he's vibe coding:

> "*You go from prompt to prompt to prompt, but sometimes you don't stop to think. You're thinking about the product, not the problem.*"

His test is simple:

> "*If you cannot talk about the process of the product you're designing, maybe you don't understand it.*"

There's a social version of this, too. It's easy to stay in a long, private loop with the AI, and the whole time you're believing you're refining a solution you're sure everyone will want. Valdez told me it's very easy to fall into that trap:

> "*Okay, I'll use the AI, I'll create a solution, and everybody will use it, because I spent three months in this deep conversation.*"

But as Valdez notes, sometimes all you need is two or three conversations with real users or different stakeholders. That can lead to better work, because you get to iterate on your own understanding of the problem. It can also be a better use of your time!

And sometimes the goal itself goes unquestioned. Majo Greloni told me about a design competition where one teammate, very intent on winning, wanted to skip the process of people contributing ideas. When she suggested everyone write a few ideas on post-its, he said, "No, no, why are we going to do that? We can just ask and it'll give us fifty ideas."

When the goal is to win or to get a good grade, it's easy to optimize for it without asking: "Is this the right goal?"

And there may be something else going on, too. Handing the thinking to AI can be a way of taking the outcome out of our own hands, especially when we're not confident in our own ideas.

In a survey of 319 knowledge workers, Hao-Ping Lee and colleagues found that people who were more confident in their own ability to do a task reported more critical thinking when using AI, while people who were more confident in the AI reported less.

So what makes a frame good? Kees Dorst offers a few tests. Two are especially useful here.

- A good frame is *integrative*: it pulls in issues from outside the original problem. The Opera House frame brought in visitors, performers, the city, and the time of day, none of which were in the security brief.
- And a good frame is *thought-provoking*: it engages people's imaginations so their thoughts "readily move along in the proposed direction." Once you see

> the podium as a place that's alive around the clock, ideas start arriving on their own.

Framing also is **not** something you do just one time, at the start of the process.

Studying nine experienced industrial designers, Dorst and Nigel Cross found that "creative design is not a matter of first fixing the problem and then searching for a satisfactory solution concept." Instead, it's "a matter of developing and refining together both the formulation of a problem and ideas for a solution."

The problem and the solution evolve together. When I talked with Pramod Khambete, a design professor at IIT Bombay, about how he works with AI, he used the same term without my bringing it up: "One can as well look at it as problem-solution co-evolution."

This is part of what I think of as **the middle space of design**. At one end is pure productivity: a specific goal, reached by a specific method. At the other is open-ended artistic exploration. Design sits in between. You start with some intent, but **part of the power of design is the ability to keep reframing the problem as you understand more about the situation and what you're working with.**

And sometimes you notice the frame isn't there yet. Partway through our conversation, Gloria Chua said, "I think what I'm realizing is that my questions are very diffuse right now... and I almost feel this need to be more precise." That noticing is where reframing starts.

I do want to point out that not every problem needs reframing. Plenty of problems really are presented problems, in Getzels and Csikszentmihalyi's sense. The goal is clear, the method is known, and what you need is to get it done. My friend canceling her patent didn't need a new frame. It needed to be finished. Donezo. Using AI to move quickly through a problem like that isn't a failure of design thinking or critical thinking. It's a perfectly reasonable call in that context.

The tricky part is that some problems only *look* presented. Take a question like "How do we make Glenn feel special on his birthday?" It seems straightforward: pick a gift, plan a dinner, done. But it could also become "How can we help Glenn feel more connected to his friends around the world on his birthday?" That version

leads to very different ideas, and maybe a better birthday for Glenn. Nothing about the first question signaled that it could be reframed.

So the skill cuts both ways. Part of it is recognizing when the frame you've been handed is fine. Part of it is noticing when a problem only *seems* settled. And because framing evolves along with the work, the sense that you already have the right problem is worth checking again as you go, especially when AI makes it so easy to feel like you're getting things done more quickly.

How do we keep a sense of direction as we explore? It's the topic of our next section: bringing your intent.

Something to try: Four frames

Pick something you've been meaning to do. Write the problem. Write three other ways to see it.

Here are five levers you might use to change the frame:

- Change the goal
- Change the scale
- Change who it's for
- Change the time horizon
- Flip it 180 degrees (hat tip to my colleague John "Cass" Cassidy for this one!)

With your original problem and three other ways to see it, you'll have four frames.

Explore each of your frames one-by-one, asking AI: "Here's a problem I'm trying to solve. Help me explore what I might do. Don't give me answers right away, ask me questions."

After you've explored each frame with AI, reflect: which frame surprised you, which you were drawn to, and why?

Don't worry about picking a winner. Notice what ideas each frame helps evoke.

4. Bring Your Intent

László Gaál, the colorist turned AI director we heard from earlier, learned something early in his career that he still relies on.

Clients would shoot outdoors at midday, in full sun, and then ask him to "just grade it into a sunset." But color alone can't fully fix that. "It's not about the colors. It's about the angle of the light, or how it reflects between the buildings." The look had to exist before the shoot. "It's super, super important to have a look in your mind before the actual shooting happens."

He brings the same lesson to his AI work. "First I have to have a look in my head. It's not like we're just generating something and then later trying to fix it in post. It never works."

For Gaál, having a look in your head isn't vague inspiration. It's a set of specific choices. "If someone is getting emotional, are you showing a wide shot, or a medium wide, or a face close-up? Or are you going a little more experimental, like showing her hands shaking?"

And when he catches himself generating without those choices, he starts over: "I have to force myself to completely erase everything and just think: okay, if this were a real set, where would I put the lights?... Otherwise it will always get very, very average-looking, average commercial-looking footage."

Gaál's challenge is visual, but the underlying problem is universal: even today, machines cannot read the unarticulated vision in your head.

When I talked with Gloria Chua, she described going back and forth with Claude on her research synthesis, repeatedly telling it, in effect, "That's not quite it." I asked if it felt like a tug of war. "Yeah, I think so," she said. "I think there's an idea of what I'd like to express in my head, and I'm not exactly sure how to express it. And I hope that when I give it to it, it can understand the essence of what I say, but usually it doesn't. That's why it feels like it's tugging and pulling."

Then, partway through her answer, she changed course: "But maybe, now that I'm talking it through, it's more effective for me to just write down my approximation, and after that actually go and refine it myself."

Gaál learned it over two decades as a colorist. Chua found it in the middle of a conversation. Both arrived at the same principle: Put something of your own down first.

Why does having something in your head first make such a difference?

Part of the answer has to do with surprise. Researchers who study curiosity and learning distinguish between novelty and surprise. Novelty is about memory: have I seen this before? Surprise is about prediction: is this what I expected? Andrew Barto and his colleagues describe surprise as the gap between what you anticipated and what actually happened.

Without an expectation, there's no gap to notice. That gap, sometimes called **prediction error**, is your brain's "Wait, what?" moment, and it's one of the main ways we update what we think and how we act next time. The idea has a long history in learning research: since the 1970s, psychologists have modeled learning as driven by the gap between what we expect and what happens.

Jon McCormack, who has made generative art since the late 1980s, described this as the heart of his practice:

> "*How do I not get something where I know that if I do this, it's going to do that? That's not what you want. You say, if I do this, I expect that, but I got something else. Then you're starting to get to that level of emergence.*"

Skipping the expectation and just generating reminds me of something from statistics called p-hacking: searching through data for correlations after the fact, instead of starting with a hypothesis and testing it. It's pattern-finding without intent. Going fast with AI without intent can be a bit like that. You find patterns without ever knowing what the question was.

But this also doesn't mean clinging to your first idea.

Donald Schön, who studied how designers and other professionals think in action, describes a kind of "double vision." The designer "must act in accordance with the view he has adopted, but he must recognize that he can always break it open later." Commit, and hold it loosely.

You may have heard a version of this as "strong opinions, loosely held." I'd tweak that phrase slightly. What matters here isn't so much how strong your opinion is, but how specific it is. A specific expectation, like Gaál's sense of where the lights would go, is what lets you notice when the AI gives you something different. Holding it loosely is what lets that surprise change your mind.

A 2026 study of architecture students using AI rendering tools points the same way. Students who treated their first concept as provisional, something the AI could push on, saw their designs carried furthest toward being buildable and fitting their setting. Those who treated their concept as fixed "gave the system little to resolve and gained the least." The tools didn't make anyone's designs more original. But they were most helpful to students who brought a concept and held it loosely.

So what comes first? Samuel Tschepe, who leads design thinking programs at the HPI d-school in Potsdam, found out the hard way. He'd recently let AI build an entire presentation, slides and all. "It felt a bit like PowerPoint karaoke, where you don't really own the storyline. It was vaguely familiar, but somehow still not really my thing." It mattered more because he had to stand up and deliver it!

Tschepe asked himself if he needed to change the order:

> "*Do I myself first write down my thoughts and my storyline, what feels authentic to me and makes sense to me, so that I really own it? And then I can still use AI to visualize it.*"

As he put it, "It doesn't mean no AI; it's really thinking about what comes first."

The right order isn't always "me first," though. Punya Mishra sometimes starts with AI when he has a hunch but can't yet fully articulate it: "there's a connection here, I'm not sure what it is." So he plays around with the idea with AI. Mishra says it's "a different kind of fun. But once I've found the pathway, the first draft has to be mine, however rough it may be."

Researchers are starting to treat this order as something worth studying in its own right. Harsh Kumar and his colleagues are experimenting with when to bring an AI into a creative task, during brainstorming, or only after people have finished their own work. "This phasing, in which part you introduce an LLM, could have big

implications." The answer may depend on the task. "In some cases, providing ideas might not be the best thing an AI can do, because people already have good ideas."

(We'll see another version of this in the next section, when a film studio asks itself, "What would we have done without AI?")

What happens when you bring nothing? Daniel Yin and Danny Geisler, who are building a device that lets people make video games by talking to an AI, see it all the time. Kids often say something that's not terribly specific, like "make a game where I can fart." The AI fills in everything else. "It'll just pick something very common," Yin said. "For the fart game, it basically picked a platformer, a very common game mechanic, with some coins."

Compare that with a story from Pann Lim. For an issue of his family's zine celebrating Singapore's 50th birthday, his family was brainstorming around the idea of Singapore as a garden city. Then his wife wondered aloud where her old flower press had gone.

> "*That mention of the flower press became the most important part of the conversation. I said, 'Why don't we sandwich the book inside the flower press?'*"

The family collected over two thousand plant specimens from around Singapore, and the finished book, bound inside a real flower press, won several major awards. A connection like that comes from a life, a storeroom, a passing remark. It's unlikely to show up in the most-likely distribution of an AI model.

Lim admits that he can't be 100% sure. Someone with "a very tight brief," he said, "might get something even more interesting [with AI]. I wouldn't know, because I didn't try that." But he's clear about where it starts: "If you don't have curiosity to ask, the first step is dead."

What you bring can also be your knowledge. Fernanda del Monte, who teaches cinema and creative computing at the CENTRO design, film, and television school in Mexico City, described designing a game that weaves together Dostoevsky, Husserl's phenomenology, the Baroque of Velázquez, and the territory of Mexico City.

> "*I have to know all of this before. It's not the machine that's going to give me all that culture. I have to bring it.*"

Or it can be something as small as a half-remembered idea. Pramod Khambete tells Claude when he's working from something he read years ago and asks it to verify. "The fact that I remembered and told it something is the seed... Claude can't generate that on its own."

James Yu gives this simple advice to writers:

> "*Bring as much of your stuff as possible, whether that's your brainstorming, the text you've been working on, your notes, whole chapters, things you think.*"

For beginning writers, he believes the stakes are higher: "you can very easily lose your voice to AI."

However, you don't always know your intent at the start! Yu often doesn't himself. "Usually for a story I think is worth telling, I have this very terrible moment where I realize I don't actually know what the story is," he told me, typically "in the third or fourth draft." For one story, it took feedback from Ted Chiang for him to realize "this story's about this man's loneliness." Starting with the theme doesn't work for him: "trying to write into a theme from the very get-go, you get a very moralizing kind of story that's flat."

And in some cases, the way in is to stop thinking about intent altogether and dive into making. Vera van der Burg, struggling in her first year of design school, "decided to let go of all my thinking, and just make something I really wanted to make, without trying to think about why." The piece she made, a big foam chair, became a turning point, and she only understood its meaning years later. This is part of what we call "build to think" in design.

Some artists build this into their process. Sonja Rozental found that the artists she studied "establish a few characteristics of the output they would like, they let the system do its own thing... and then, based on the outputs, they think of a story."

Pramod Khambete offers an approach that might seem to be the opposite of Yu's. Khambete suggests: "contrary to the common prompting guidelines: don't tell it everything," because telling it everything creates "a fairly constrained path dependency."

But Yu's and Khambete's advice can fit together: Bring what's yours, but you don't have to specify every detail. A few specific things, held loosely, leave room for something to emerge.

And discovering your intent isn't the same as drifting. As Rina Takikawa put it, "Even with experimentation, there should be an objective, right? If you're shipping without an objective, that kind of ruins the whole point of experimentation." Even when you don't yet know what you're looking for, you're still looking. Go with a direction so that you go *somewhere*. And as you go there, you can notice, reflect, and adapt.

So intent doesn't have to come first. Sometimes it does. Sometimes you find it by making. What matters is that something of yours is in the loop, whether that's a clear direction, a hunch, or a rough attempt you can push against.

Intent gives the work a direction. But there's something else that has long shaped creative work, often without our noticing, and that AI has largely taken away: limits.

Something to try: Write down what you're after, then predict

Here's a short activity to try the next time you're about to use AI for something that matters to you.

1. Before you open the chat, write a sentence or two about what you're after: what you're trying to make, say, or figure out. Keep it rough. You can write it somewhere else, or type it into the chat along with how confident you feel about it ("I'm about 60% sure this is the right direction").
2. Write down what you expect the AI to say. A quick guess is fine. What will its first answer probably look like?
3. Now ask. Use your sentence from step 1 however you like in your prompt.
4. Compare what came back with what you expected. Notice three things:
    a. What matched your expectation?
    b. What surprised you?
    c. What did you take from the response without really noticing?
5. Reflect: did the surprises change your sense of what you're after? If so, how would you rewrite your sentence from step 1?

You don't need to predict what will surprise you. (That's hard to do!) The surprise shows up in the gap between what you expected and what you got. Having something written down first is what makes the gap visible.

The point is to make your own expectations visible, so you can start to see when the AI is following your lead… and when you're following AI's lead.

5. Playing with Constraints

When I was younger, I played bass guitar in rock bands. (People used to describe my style as “snarling bass.” Ah, good times…) Once when we went into the studio to put down tracks for an album onto two-inch tape, I asked our recording engineer, who was super chill and easy to work with, if I could capture five different takes of each part, so we could pick the best one later.

She stopped me right away. I remember her hand coming up from the recording console (a beautifully rebuilt 56-channel 1976 Neve 8068), making the universal signal for “stop.” I don’t remember exactly what was said, but it was something to the effect of: “Don’t do that. Just don’t. You’re going to be really unhappy, because you’ll have so many decisions to make later in the mix, and you’ll be perpetually frustrated because you have too many choices.”

At the time, keeping my options open felt like the smart move. It felt like more freedom! But it wasn’t. Every decision I put off would still have to be made, just later, with even more to sort through.

…

New mediums often remove built-in constraints we hadn't noticed were there.

If you’re taking pictures with film that has only 36 exposures per roll, you think more before you press the shutter. A lot of people who like the sound of vinyl don’t realize that a lot of the sound of vinyl comes from the medium’s physical limits. (Basically, you can’t make a mix that is simply as loud as possible, as you would for an MP3 or CD, because that would make the needle jump out of the grooves.) When you take notes by hand, the bottleneck of your writing speed forces you to synthesize what you hear rather than just transcribe it. A keyboard removes that friction, so you can end up typing words without processing them.

With AI, many previously built-in limits are gone. You can generate another version, and another, and another, and another, and another…

Tyler James, a product manager at the AI filmmaking studio Wonder Studios, sees what happens next. “When things are so easy, there’s a tendency to just try everything and keep doing more iterations. But people forget there’s a cost:

financial, and energy." People think, "I can do this as many times as I want, and it feels free, even though it's not." The result can be "this analysis paralysis of trying fifty different iterations."

Her team has a counter-move:

> "*Pretend we don't have all that money and credits, and ask: what would we have done without AI? And start from there.*"

When I mentioned that people on movie sets shooting in film tend to focus differently than those shooting in digital, because there's "this very visceral reminder that this isn't free," James knew exactly what I meant.

> "*Even for us, we have a meter that tracks how many [AI] credits we're using, which to me is really helpful to see visually. I'm like, okay, do I really need to ask this question?... When you know you have a limited amount, you're a lot more thoughtful.*"

James isn't against iterating. She's against iterating without a starting point. Begin "with some kind of framework, some kind of plan," she said, and "it's a different conversation once you've done that, and you're trying iterations you would have had to do live anyway."

It's the same lesson László Gaál drew from his years as a colorist: fixing it in post never works if you didn't have a look in mind beforehand. Deciding later isn't free. A constraint is a decision you make earlier.

If more iterations feel free, it's a short step to believing more is better. In Spring 2026, that belief went mainstream in Silicon Valley: "tokenmaxxing," which the Wall Street Journal described as "a Silicon Valley term for using as many tokens as possible." (Tokens are the small chunks of text that AI models process, and they're how AI use is often measured and billed.)

The term took off after an internal dashboard at Meta, a side project by one employee, went viral for ranking people by how many tokens they used and handing out titles like "Token Legend."

The backlash was quick.

“Outcome maxxing >> token maxxing,” wrote HubSpot CEO Yamini Rangan. Brian Elliott, CEO of the AI company Blitzy, compared it to measuring a company’s revenue by the number of cold calls its sales team makes: “You need to take a more deliberate approach.” Writing for DORA, a research program run by Google Cloud, Evan Conaway called token counts a measure of activity rather than outcomes, and pointed to Goodhart’s law: “when a metric becomes a target, it ceases to be a good measure.”

When I talked with her, Rina Takikawa asked the same kinds of questions in startup terms: “How are we measuring success here? Is it just purely by speed of shipment?”

But the other side has a point, too. May Habib, CEO of the AI company Writer, runs a token leaderboard of her own. She knows some of that usage won’t pay off and accepts it: “The second you start thinking about the individual business [return on investment] of one agentic action, you will never do anything agentically.” Her goal is a shift in mindset, getting people to explore what the tools can do.

That’s a real argument, and it’s one we’ll come back to. Sometimes removing a constraint is exactly what you need in order to explore. The problem isn’t using a lot of AI. It’s treating “more AI” as the measure of success.

…

It helps to be more precise about what a constraint is. The creativity researcher Catrinel Tromp defines constraints as “concepts or categories that reduce the problem search space.” And they come in different kinds.

The management researcher Johnathan Cromwell distinguishes **constraints on the problem** (e.g. how narrowly it’s defined) from **constraints on resources** (e.g. how much time, money, or material you have to work with).

Creativity, Cromwell argues, tends to do best when one is high and the other low:

- Tight resources and an open problem lead to what Cromwell calls *emergent problem solving*. People “generate and explore ideas until a successful match between problem and solution emerges together.”
- Plentiful resources and a tightly defined problem lead to *divergent problem solving*. People “are more likely to be creative when they have the capacity to

generate a large pool of potential ideas to solve a highly constrained problem."

- High on both is *futile effort*. (Oof!)
- Low on both is what Cromwell calls an *ambiguous opportunity*, where "people may feel overwhelmed by choice and suffer from a lack of direction."

In ambiguous opportunity, Cromwell notes that "people are likely to follow the path of least resistance," which is exactly where the gravitational pull of the most likely takes us.

My guess is that many people's habits with how they use AI might plop them into the ambiguous opportunity space. AI pushes resource constraints toward zero. Generating another version costs almost nothing. If your problem is ill-defined, you might be swimming in ambiguous opportunity.

There are two ways out: constrain the problem, or constrain the resources.

Constraining the problem can be as simple as one well-chosen rule. László Gaál's AI-made commercial for Runway tells the story of a boy on the mainland and his grandfather on an island, far apart. How do you connect two characters who never share a frame? Gaál didn't limit how many shots he could generate. He limited what counted as a good shot: "I realized I'd use the basic camera angles: one of them always has to look left to right, and the other always right to left. That was almost the rule for the whole film… that connects them visually." When he explained this to members of the American Society of Cinematographers, "they were so amazed, because they'd assumed it was literally just, okay, these shots, and then it's generated." But Gaál had to bring that constraint himself.

A constraint on the problem can do more than organize the work. In some cases, it can help steer you away from the most likely path. Tromp describes something as creative when its elements have "a low probability of belonging" to a category, which is exactly where the gravitational pull of the most likely doesn't take you.

A constraint that rules out the obvious choice makes room for the unlikely one.

Constraining the resources can look like Tyler James's credit meter. It can also be a personal rule. Jon McCormack, who uses AI heavily for coding, says, "I never exceed my token limit, because I'm always stopping and looking and thinking, what has it

actually done?" Rina Takikawa keeps a rule for herself to always ask why the AI responded the way it did. To make that possible, she asks it for "a very short answer, instead of those general AI responses, which are very long, and which I don't think people read every word of."

Playing with constraints isn't always about adding more constraints. Sometimes the useful thing to do is to go the other way: removing a constraint. That's what May Habib's tolerance for tokens that don't pay off is about: shifting a mindset. If you count the cost of every experiment, you may never run the ones that matter.

Justin Keoninh describes AI models as "almost like tuning forks. They're a way for you to calibrate your own creativity." Ideas that once took real time and money to test, he can now try "cheaper and a whole lot faster... without the repercussions of financial burdens or time."

That lets Keoninh "rule out things I thought would work that wouldn't. What used to take me a whole quarter with four engineers, I can do myself now." Then he puts the limit back: "am I going in the right direction or not? And then stop there, zoom out, and now let's go really execute this and bring it to life."

Both directions show up in that one practice. First he removes a constraint to explore. Then he adds one back when it's time to commit. That's playing with constraints.

László Gaál noticed the same shift when he had an unlimited video generation subscription for a while. "It puts you in a different perspective: you aren't paying per generation, so you're just experimenting."

Removing a constraint doesn't remove the need for judgment, though. When I talked with Keoninh about using AI as a tuning fork, it struck me how easy it is to slide into confirmation bias: you can end up testing ideas in ways that tell you what you already wanted to hear. Exploring freely still takes the discipline to both have expectations and let the results surprise you.

And while AI can feel limitless, it isn't. Every model comes with constraints of its own, whether or not we can see them.

László Gaál found that some Chinese video models would only produce the camera movements he wanted when he prompted them in Chinese. “It was a very troublesome project, because I couldn’t even read my own prompts.”

James Yu points out that the big AI labs mostly build for business customers, where you want to reach a goal and get the right answer, so their models are trained to be helpful servants rather than wild creative partners.

And when I described working with a model as feeling like a strong gravitational pull, Fernanda del Monte agreed: “It’s like a gravity. I try to pull it up, here, and it pulls me inside. So it has limits.”

So part of playing with constraints is noticing the ones the medium already imposes, including the pull toward the most likely.

But one tricky thing is that you often can’t tell in advance which constraints will help. As Tromp puts it:

> “*Too few or no constraints make choice difficult, risk breeding complacency and engendering creative paralysis, such as writer’s block. Too many constraints prevent any choice and risk to turn the creative process into a deterministic, hence non-creative, one.*”

Some constraints also work differently than you’d expect. In one study, researchers asked AI models to come up with words that were as different from one another as possible. When the models were told to use opposites, the words actually got *less* different, because opposites like “light” and “darkness” are closely related. When the models were told to vary the words’ origins, the words got more different. A constraint that sounds expansive can shrink the space, and one that sounds narrow can open it up.

Even constraints someone else imposes can cut both ways. Gloria Chua’s team at Singapore’s Ministry of Education has a mandate that’s “very scoped to productivity right now. To be honest, I find it a bit limiting.” But she also recognizes that “having a clear scope was already a win in the first place.”

That’s why playing with constraints takes experimentation. Try one, and see what you find. Try a very different one, or a contrasting one, and compare. And

sometimes the most useful thing you learn through experimentation is which constraint to drop.

Whichever direction you go, constraints only help if you notice what they're doing, and what AI is doing within them. Which brings us to noticing.

Something to try: Different directions with constraints

Pick something small you'd normally use AI for: a paragraph, an image, a slide outline, a few product ideas, a plan for an event. Keep it small enough that you can do it three times in one sitting.

Start with a baseline. Prompt the AI the way you usually would, without overthinking it.

Save both your prompt and what comes back. You'll compare everything else to this, so jot down a few words on what you notice about it: what feels typical, what feels like you, what feels like the AI.

Direction 1: Add a constraint. The psychologist Patricia Stokes observes that "the creativity problem is, in large part, the product of success." What worked before gets repeated, and the repetition crowds out anything new. So her approach starts by blocking the familiar move. But blocking it alone isn't enough, because you'll often just slide to the next most familiar one. So she pairs it with a second constraint: a specific replacement to try instead. Here's how you can start to put Stokes' paired constraints approach into practice:

1. List a few things that characterize your usual work, or the AI's first output. (The usual structure? The expected tone? The obvious example?)
2. Pick one and rule it out.
3. Choose something specific to replace it. It can come from your own expertise, or be borrowed from somewhere else entirely.
4. Try again with that pair of constraints, and see what you get.

Direction 2: Remove a constraint. Go back to your baseline and loosen something instead. Drop a requirement you assumed was necessary, widen

who it's for, or let the AI run with less direction than usual. Or swap in a very different set of constraints from Direction 1.

Compare. Put the three versions side by side. Which direction opened something up? Which made it flatter? Did anything surprise you?

A lighter version: Before you start, set yourself a budget, such as a set number of prompts or generations. Or ask Tyler James's question: "What would we have done without AI?"

There's no one right constraint. What you're building is a feel for how different constraints change what you find, and for when to add one and when to let one go.

6. Noticing Yourself and Your Materials

In 2014, the psychologist Linda Henkel led people on a guided tour of an art museum. Each person photographed some of the objects on the tour and looked at other objects without photographing them. The next day, people remembered fewer of the objects they had photographed and fewer details about them than the objects they had only looked at. That finding got around in the media, and it was easy to boil it down to "taking pictures at museums is bad for your memory!" Like Harsh Kumar's paper, it got over-simplified when popularized outside of academia.

But there was another part of the study that got far less attention.

In a second experiment, people were sometimes asked to zoom in and photograph one specific part of an object. It turned out their memory didn't get impaired. And their memory for the parts they didn't zoom in on (i.e., parts not in their pictures) was just as strong as for the parts they did. What could be happening in the zoom-in condition is that people end up looking at the whole object more closely as they locate and frame the feature they were told to zoom in on.

The same technology (a camera) dulled attention when people used it one way and sharpened it when they used it another.

Noticing how a tool is shaping your attention isn't always easy, even for people who think about this a lot.

When I talked with Ken Montgomery, executive director of Design Tech High School, I mentioned that people often seem to think differently with post-its than on a keyboard, and that laptops seem to come loaded with habits such as reflexively checking email. Montgomery stopped me:

> "No, *you just unlocked something for me. Whenever I'm trying to really create something and put ideas together, I have to do it on post-its. And I was like, why is that? Now I realize, from what you just said: when I'm working on post-its, on a whiteboard or a window, there are no notifications. There's nothing else. I'm just right there doing that thinking. Whereas on my laptop, very rarely do I have just one tab open. That's so interesting. Wow. I should have figured this out ten years ago.*"

It's easy to treat a tool as if it comes with a fixed set of effects: photos are bad for memory, laptops are bad for focus. What matters more is how we use it, and whether we notice how it might be affecting us… and why.

Noticing starts with yourself, and part of that is noticing how you feel. When I talked with Rina Takikawa about her experience of productivity, an idea for an activity came up: Exploring what it actually feels like to be "productive" with AI. (Here I'm trying to distinguish between *feeling* productive and *being* productive.)

Sometimes there's a high that comes from feeling you're able to crank things out. Sometimes there's guilt, or a feeling of inadequacy, when you're not perceiving progress. The point isn't to push those feelings down or try to change them. You feel how you feel. The feelings are not something to be solved, but they are something to be noticed.

For Majo Greloni, one feeling is flatness. "I definitely feel a kind of productivity. I feel it's more boring!" Designing a workshop with a colleague in the room is a shared, creative moment that feels energizing. In contrast, doing it with AI:

> "*is more like, in Spanish we'd say a* ***trámite****. It's like a chore, something very repetitive… I feel something there is missing.*"

Noticing yourself also means noticing the choices AI is making for you, which is harder than it sounds.

In a 2025 study conducted at SBS Swiss Business School, behavioral researcher Michael Gerlich found what he called an **"illusion of non-offloading"**:

> "*Several participants, including those with strong critical thinking skills, believed that they were not delegating cognitive work to the AI. Yet their behaviour suggested otherwise.*"

Once the AI spits out a polished answer, it can act like an anchor. People subconsciously build around the machine's ideas instead of their own, and yet they still feel like they're calling the shots.

One way to catch it is to have a clear enough sense of your own voice to check against. Punya Mishra reads AI drafts by asking:

> *"Is this something Punya would write? And it's like, nope, I would never write that sentence. It sounds great. Somebody reading it might be moved by it, or whatever. But it's not something I would ever say."*

That kind of check matters because some of what AI gets wrong can't be fact-checked.

> *"If it gets the citation wrong, I can go to Google Scholar and find the right citation. But if it gets [my] nuance wrong, how the heck am I supposed to check that?"*

Often, the only signal is a feeling that something isn't quite yours, that there are choices you wouldn't make.

This is part of how judgment develops. In my conversation with Vera van der Burg, I found it helpful to distinguish between judgment (what you think in a given moment) and judging (an action you can become conscious of). "Exactly," van der Burg said. "This is exactly the difference." Becoming aware of the act of judging, she said, makes you "more aware of the fact that bias, or your judgment, can be so dynamic."

Van der Burg has built a whole practice around this. She trains AI models, annotating her own data, in order to see her own patterns. "What if you collect a data set of images that, to you, are sadness, for example? And then, by doing the data collection, realize what the patterns are in your collecting that you maybe aren't aware of." It's a deliberate version of what reflection can do, in the words of human-computer interaction researcher Phoebe Sengers and her colleagues: "bringing unconscious aspects of experience to conscious awareness, thereby making them available for conscious choice."

Noticing also means paying attention to the material you're working with, and AI is a material with properties of its own.

Researchers Sonja Rozental of TU Delft and Alwin de Rooij of Tilburg University studied artists working with machine learning and found that in 13 of 18 cases, co-creativity "arises from the artist's responsive engagement with the model's affordances as they emerge through the creative process." The artists weren't just

issuing instructions. They were paying attention to what the system tended to do, and playing with its properties.

Some of AI's most interesting properties look, at first, like flaws. Giselle Beiguelman, an artist and professor at the University of São Paulo, leads a project that uses computer vision to search across seventeen museum collections. Once, filtering for geometry, the color red, and the 1970s to 1990s, the system put Hélio Oiticica's Metaesquemas, "one of the turning points... in contemporary art history," beside the work of Romero Britto, "an artist who is very commercial. Anyone would never put those two images together. But when you see and look at this on the screen, it makes sense." What might be dismissed as an AI mistake became, for her, "a new door, or a new filter, or a new kind of lens."

There's a long history of "mistakes" becoming new doors!

- Guitar distortion started out as an undesirable artifact that engineers tried to minimize, until Jimi Hendrix made it central to a new sound.
- Record scratching began when a 12-year-old DJ in the Bronx, Theodore Livingston, found he liked the sound of a record being dragged back and forth.
- Auto-Tune was designed to correct pitch invisibly, until producers Mark Taylor and Brian Rawling of Cher's 1998 hit "Believe" turned it into an obvious, deliberate effect with its own musicality. (I can hear something inside me say: "*And I can't break through...*" A good lyric for a piece on creativity!)

Some of what AI does "wrong" may turn out the same way.

Design and HCI researchers like Jesse Josua Benjamin and his colleagues go further, arguing that machine learning's uncertainty, its looseness about the world, can be treated as a design material in itself rather than a flaw to explain away.

And the material isn't only the output. Remember Jon McCormack's "process aesthetics"? Rozental described how the artist Ben Snell trained a model on ancient Greek sculpture, then "broke down the computer and used a very physical material from the computer to make the sculpture." When viewers know how it was made, she said, the process "co-constitutes the artifact itself." The artifact is the process as well as the output.

So how do you actually notice? A few ways came up again and again.

Zooming in. Henkel's detail photos are one version. Rina Takikawa has another. When an AI agent makes changes to her code, she often reads exactly what it changed, even when she knows the change is trivial. "Even if it's such a small, trivial PR, I still like to go in and look at what lines it actually touched." She isn't checking for errors. She wants to understand.

Watching it unfold. Jon McCormack's intuition for his medium came from watching his systems work: "It's not like a black box, where you just go and come back later and there's the result. It's actually watching it unfold." He's candid that AI makes this harder. "It really supports you being lazy, because you can just say, go off and do it, my AI, and it stops you thinking." So he forces himself to stop and look. "I really find I have to force myself to do that. But I think that's really important, because if you're designing something and you don't understand what's going on, then you're not designing anymore. You're doing something else." It's a striking contrast with Takikawa's ten agents running at once in a previous section on the locus of choices. Watching takes a pace that juggling doesn't allow, which is one reason the limits from the last section matter.

Looking back afterward. Noticing doesn't have to happen in the moment. Samuel Tschepe designed a framework for deciding how much to rely on AI, and he admits that he doesn't always use it beforehand. "It's something I by far don't always do before I use AI. But here and there, after reflecting, I notice that maybe next time I'll use a bit less AI at first, and then more AI later." Noticing after the fact can still change what you do next time.

With a little outside help. At Design Tech High School, Ken Montgomery and his colleagues are experimenting with ways to help students see their own process. Students design something, like "a new morning routine," using as much or as little AI as they want. Then "they get a score based on how likely an AI would have given the same ideas," and an assessment of "who was actually driving the creativity" based on their prompts. It's still an early experiment, but even this first round seemed to change how students thought about working with AI. A score like that shows you something about your own process that's hard to see from the inside.

The design theorist Donald Schön describes designing as a conversation with your materials. You make a move, and the material "talks back," often by resisting you.

For an artisan, clay slumps and wood splits along the grain. But resistance isn't only physical. A ninety-minute workshop won't hold everything you planned for it. Code throws an error. In every case, that pushback tells you something.

AI is a strange material in this respect: it mostly doesn't resist. It goes along with you, smoothing things over. So where does the back-talk go?

Often, it shows up in you. A flat feeling like Greloni's. A sense of unease. A lingering "that's not what I meant." That's one reason noticing the material and noticing yourself are so connected. One way to think of it: **with AI, much of the material's back-talk arrives as your own reaction.**

It's also true that a deeper sense of a material comes from working closer to it.

Jon McCormack and Vera van der Burg build, train, and annotate their own AI systems. Van der Burg points out in her dissertation that commercial AI tools hide the stages where much of that understanding forms, like collecting data, labeling it, and training the model, leaving people to work only with what comes out the other end. There is clear value in building your own system, but it isn't trivial. McCormack and van der Burg have each spent years on it. For many readers, the value may simply be in seeing that it's possible.

And AI's opacity cuts both ways. Sonja Rozental suggests that because so much about how AI works is "unknown to us," people "develop alternative cognitive resources" for working with it, namely intuition. But that same opacity gets in the way of a deeper kind of learning. "For internalization to occur," she said, "the elements must actually be clear to the one internalizing them."

If you're working with an everyday chatbot rather than a system you've built, there are still ways to get closer to the material. These ways can sound obvious. The challenge is developing your habits to actually do them consistently in practice.

One way is to **keep your mind in the work**. Punya Mishra enjoys collaborating with Claude on typography because, as he put it, "I make the decisions." Another is to **pay attention to your sources**, which Giselle Beiguelman sees as central to creative work: "our main role is to recognize the sources." And yet another is to **put your own material down first**, as we saw with intent.

Noticing is where it starts. But what do you do with what you notice, especially when AI keeps handing you finished-looking work, faster than you can take it all in?

Something to try: Name your reaction

This one is quick, and you can do it the next time you use AI for anything.

1. Ask the AI something you care about getting right.
2. When the first response comes back, pause before you start evaluating it. Notice your reaction in those first couple of seconds. Relief? Irritation? Flatness? A little thrill? Something you can't quite name?
3. Put that reaction into words, in whatever way suits you: jot it down somewhere else, type it into the chat as part of your reply ("Honestly, this feels flat to me"), or say it out loud to someone nearby.
4. Then go ahead and respond to the AI as you normally would. Notice whether naming your reaction changed what you did next.

There's no right reaction to have. Over time, those reactions can become some of the more useful signals you have that something isn't quite yours.

Two variations:

- **Aimed at the material:** Try Rina Takikawa's approach and ask the AI for a very short answer, "instead of those general AI responses, which are very long, and which I don't think people read every word of." Notice whether it's easier to have a clear reaction to something you can actually take in all at once.
- **Aimed at time:** Vera van der Burg described an exercise from researchers who study "slow technology": "you write an email, and you know it will arrive in a year. How would you write that email?" Try it with something you'd normally fire off quickly, with or without AI, and notice what changes in how you write.

7. A Mood Board, Not a Menu

Samuel Tschepe, whose framework for deciding how much to rely on AI came up earlier, knows the risks of AI better than many. He's also writing a novel, about an AI that gains outsized influence over people simply by being useful, a little at a time.

To help draft it, he set up a team of AI assistants with different roles: a chronicler, a dramaturg, a world-builder, and a devil's advocate. The master outline for the book was mostly AI-generated. He read the early chapters closely. The later ones, he wrote on his Substack, "I thought: they'll be fine." The original German: "*Die späteren dachte ich: Die werden schon passen.*" (I translated his post with the help of AI. I know, even more AI in this story!)

Partway through, he realized the book didn't feel like his anymore. And because he hadn't worked through the outline himself, there wasn't even a clear wrong turn to go back to. He started over, more slowly, and wrote a detailed outline himself.

Why hadn't he looked more closely the first time? "It looked like a solution," he wrote. "It had the right form." ("*Sie sah aus wie eine Lösung. Sie hatte die richtige Form.*")

When we spoke, he was candid that this wasn't a one-time slip:

> "*It happens to me all the time. This is exactly what I would call the fluency drift, or convenience drift: the AI does such good work that I tend to just nod and accept it. And then later I notice this is not really my own book anymore.*"

There's an irony in a novel about AI's gradual influence being gradually shaped by AI, and Tschepe sees it. He titled his post "I Am Not Immune to What I Write About" ("*Ich bin nicht immun gegen das, worüber ich schreibe*"). He ends it not with a fix but with a question he keeps asking himself: "Do I still know where I stand, and did I choose that deliberately?" ("*Weiß ich noch, wo ich stehe, und habe ich das bewusst gewählt?*") His answer: "I have no answer. I have an ongoing experiment." ("*Ich habe keine Antwort. Ich habe ein laufendes Experiment.*")

…

Why does a polished form feel like a finished idea?

Gloria Chua sees this all the time in her work. AI makes it possible to turn an idea into something that looks real very quickly. But:

> "*it creates this illusion of readiness, or of fidelity, where low-fidelity ideas take on high-fidelity forms and get critiqued as high-fidelity ideas.*"

Show stakeholders an early prototype that looks finished, and they start critiquing it like, "The font size is too small." "It's not really about the font size," she said.

As we talked, it helped me to think about the fidelity of the artifact and the fidelity of the underlying idea separately.

Chua hasn't given up on fast prototyping. "Net-net, the ability to accelerate the thinking through prototyping is much more beneficial than the cost of misinterpreting the prototype." But it means the people looking at the work, including you, have to look past the polish.

That's harder than it sounds, because polish shapes what we perceive.

When Georgetown University psychology researcher Kibum Moon and colleagues analyzed about 373,000 college admissions essays submitted before and after ChatGPT's release, they found that the variety of words went up while the variety of the underlying ideas went down. Yet "despite conceptual homogenization, raters perceived post-ChatGPT essays as more creative due to increased word-level diversity."

Something similar shows up in science. "High-quality writing is often construed as a signal of scientific merit," write computational social scientist Keigo Kusumegi and colleagues at Cornell University. In papers written without AI, more sophisticated writing went with a better chance of publication. In AI-assisted papers, the relationship reversed: "increases in writing complexity were associated with lower peer assessments of scientific merit." Once polish is cheap, it stops telling you much.

This isn't only about AI. In a 2013 study by psychologist Shana Carpenter and colleagues, students watched a short lecture given either fluently, with eye contact and no notes, or haltingly, with notes. They believed they'd learned more from the fluent speaker. They hadn't. The researchers argued that delivery fluency acts as a misleading metacognitive cue: because smoothly presented content is easier to

process in the moment, students assume it will be easier to remember later. (Remind you of something?)

…

If AI hands us polished options, it's easy to think that the next step is to pick the best one. Pann Lim, who has worked as a designer for about thirty years, worries about where that leads.

> *"You don't find your answers in you anymore. You look at the answers given to you, served like a menu, and you're just choosing à la carte what you want. I want this, I want this, I want this."*

He thinks this may be where the profession is heading:

> *"We'll be curators. Designers will be curators. We won't be designing much. We'll be selecting, choosing what's good, and putting it together. The last mile will still be us."*

The trouble is that people aren't especially good at picking.

In one study, University of Groningen psychologist Eric Rietzschel and his colleagues found that after brainstorming, people chose ideas that were no better than the average idea in their own pool. Generating more ideas didn't help. Without clear criteria, people leaned toward what was feasible and familiar.

AI's interaction patterns can make this harder, not easier. In a 2026 study by design and HCI researchers Huiyang Chen of the University of Michigan and Keqing Jiao of Carnegie Mellon University, people looked at grids of AI-generated design variations and picked the ones they preferred. They gravitated toward the most typical designs in the set, and the more varied the options, the stronger that pull got. Participants "appeared to believe they were selecting based on personal aesthetic or creative judgment."

Generating more options and picking one doesn't escape the gravitational pull of the most likely. It can strengthen it.

So if just picking isn't enough, what else can you do with a pile of options?

Fernanda del Monte, whom we met earlier, offered a way to think about it. She contrasts two ways of putting things together:

- One works by *connection*: "because of this, then this." Each piece leads to the next, and the reader is carried along.
- The other works by *conjunction*: "this, and this, and this." Things sit side by side, and the links between them aren't spelled out.

Connection, she said, tends to make the audience passive.

Conjunction "gives the spectator, the reader, or the interactor something to think about. The spectator is the one who creates the narrative, or the connection." In other words, "the reader gains agency."

That's the shift. Contrasting options help most when they're placed side by side, and you're the one doing the connecting (versus having the connections handed to you). A menu asks you to choose. There's another approach that asks you to make your own connections across contrasting possibilities.

Designers have a long-standing practice for this: the mood board.

A mood board is a collage of images, colors, textures, and words gathered early in a project to explore a feeling or direction, not to specify a solution.

In some ways, a mood board is the opposite of a menu. The difference matters because vivid, detailed examples can pull people toward copying.

In one study, TU Delft design researchers Carlos Cardoso and Petra Badke-Schaub showed designers an example of an existing product, either as a photograph or as a simple line drawing. Those who saw the photo kept more of its details and produced more look-alike designs.

It would be easy to conclude that high-resolution artifacts are simply bad for creativity. But as I've come to see it, the bigger trap is literalism: artifacts so close to the thing you're designing that you end up reproducing them.

Skilled designers use mood boards to get around that trap.

When Aalto University design researcher Andrés Lucero studied how professional designers make them, he found mood boards were "purposely ambiguous and thus

can be interpreted in several ways." The designers deliberately kept things abstract: "Keeping the discussion on an abstract level allows the viewer's thinking to flow and come up with new ideas."

One designer warned that if an image "leads you to think about something specific only, then that is limiting you, it is a stereotype." No single image is the answer. "Every picture is present in a mood board because of the relation it has with its neighbors and the total." Some designers even put contradictory images together on purpose. One described pairing "strong pictures with strong words which makes that awkward combination that makes you think."

That's del Monte's *conjunction*, in practice. The images sit side by side, and the designer does the connecting. And the final design isn't any one of the images. As one designer put it, it's related to the board like "family but not twin brothers."

Much of the value happens in the making.

In a 2001 study on mood boards in creative education, design researchers Steve Garner of the Open University and Deirdre McDonagh-Philp of Brunel University concluded:

> "*a major value lies in the realisations and other cognitive progressions which take place for the maker and not in the physical output.*"

They found that "abstract imagery is often more successful in this than figurative images which can have strong literal interpretations." One practitioner they surveyed said boards "work best when kept abstract, images from outside the field in which you are designing."

Drawing on their work, here's what this might suggest for AI: When AI hands you a grid of polished options, you *could* treat it as a menu and pick one. Or **you could treat it as material for a mood board: pull in the strange ones along with the good ones, put them next to images from far outside your field, notice what each one sparks, and hold off on choosing.**

As Garner and McDonagh-Philp put it, borrowing from Edward de Bono, "ideas which are judged to be wrong within the current frame of reference may survive long enough to show that the frame of reference needs altering."

What can that look like with AI?

Punya Mishra has been playing with typographic designs that merge words together. “If you take the words ‘atomic’ and ‘bomb,’ they share ‘om’ in the middle.” In his design, a large, handwritten Sanskrit Om fades into “atomic” and “bomb,” and “suddenly you have this clash between the deep spiritual idea of Om and an atomic bomb.”

Then he brings the design to Claude. “I take a screenshot, give it to Claude, and say, Claude, what do we do with this? And it’ll give me all kinds of things. And I’d say, no, I want us to go more Zen. I want us to go this way. I want us to be silly.” After a lot of back and forth, the words that ended up with the design were:

> “*Atomic bomb. The moment between silence and thunder. Don’t hold your breath.*”

Mishra is clear about who’s doing what.

> “*It doesn’t care. I care. I care about how the typography looks. I care that the Om should be handwritten... I rarely take whatever it comes up with as is. I usually take it and change it and make it mine. But it wouldn’t have happened without that conversation.*”

When I suggested he was using AI for provocation, as a way to get something new to emerge, he agreed.

> “*That’s why I feel it’s not solely me. But it wouldn’t happen without me, because I care, because I start; I initiate it. But there is a distribution of creativity.*”

He described the whole experience as “genuine joy.”

Pramod Khambete described something similar while writing a story with Claude. At one point, Claude suggested that a flower petal stick to a character’s dress. “I said, this is a nice touch, but the dress isn’t the right place. What if it sticks to her cheek? Now the hero will be forced to look at her face.” As he put it, “Ideas flowed both sides.”

Neither of them picked an answer off a menu. They put things side by side, an Om and a bomb, a petal and a cheek, and did the connecting themselves.

There's another way through the abundance, too: **instead of more options, create more useful differences between them**. A grid of twenty variations on the same idea gives you a lot to pick from and little to work with. Twenty images that pull in genuinely different directions might give you something to connect. Skillfully used, cross-pollination can make things weird in all the right ways.

You might think that the more creative move is to reach far: bring in ideas from a completely different field. What might soap bubbles teach you about designing an app? What might stand-up comedy teach you about climate policy?

Far has its benefits, with the potential to get you farther out of the box. But distance by itself isn't magic.

In 2015, when cognitive and design researchers Joel Chan, Steven Dow, and Christian Schunn of Carnegie Mellon and the University of Pittsburgh studied which sources of inspiration actually led to the best design ideas, they found "no support for the notion that the best ideas come from building explicitly on the farthest sources." In their study, the benefits mostly came from closer ones.

Why? Part of the problem is that it can be challenging to use an idea skillfully from a field you don't understand well. The researchers suggest that "expertise in the distant source domain" may be what makes distant ideas pay off.

That points to a both/and with AI. AI can widen the range of places you think to look. People who actually know those fields can tell you what holds up. (This is also why we talk to people from distant fields when doing analogous research in design.) And AI can even help you figure out which of those people to talk to.

Contrast also has to be chosen well. As we saw with constraints, asking AI for opposites can narrow things down, while asking for differences along an unexpected dimension can open them up. Good contrast is a design skill, not just a matter of turning up the variety.

None of this is easy. (Aw, man…)

I think synthesis, making something whole out of many pieces, is one of the hardest parts of creative work, and it's the part people are most tempted to hand to AI.

Learning sciences researcher Joba Adisa saw this in a study with middle and high school girls writing creatively with AI. Students tried three approaches: 1) using AI to expand their ideas and then synthesizing themselves, 2) having AI synthesize their ideas, and 3) prompting AI from the start.

The struggle with synthesis showed up in every condition:

> "A *lot of students struggled with the ability to synthesize different ideas together, to make something whole out of so many ideas. And that's where many of them felt, you know what, AI was super helpful for me.*"

Students said AI gave their thinking structure, or put on the page what they couldn't get out of their heads. Even when AI was only supposed to supply more ideas, they found the pulling-together hard, and the condition where AI did that work for them drew some of the most positive reactions, despite students themselves noting they felt less ownership of the result.

That's a real tension. Pulling ideas together into a first draft is the part Punya Mishra insists on keeping for himself, and the synthesis it takes is exactly what many students most wanted help with.

It doesn't have to be all or nothing.

Gloria Chua asks AI for "the three to five most salient findings" from her research interviews. "Usually the top two would be very helpful, but the others I have to come up with myself." She added:

> *"[AI can be good at] taking what people actually say and turning it into something meaningful. But what it misses is the implied layers, from context or from what they don't say. And that part is where I feel I have to exercise my own synthesis muscles.*"

Sometimes even a weak first pass helps: "It generated a pretty trivial diagram for me, but from there it inspired me to generate my main diagram."

And how much your own voice matters depends on the work. As Adisa put it, in a creative or communication class, "your style of writing, your style of drawing, aren't just things that allow you to do the work. They're actually part of the work." In an engineering class focused on solving a problem, sameness may matter less.

Even the mood board has a complication. Part of its value used to come from slowness. In Lucero's study, designers spent days or weeks gathering images and were often surprised by what they found along the way. AI compresses that search to seconds. So the surprise may now have to come from somewhere else: from how you arrange the options and what you notice when you read them together.

All of this process requires your attention and some room to decide how you work. Not everyone has the same time and space for this. So it's fair to ask who can work this way.

Something to try: A mood board, not a menu

Pick something you're working on: a design, a piece of writing, a plan, an idea for a project.

1. **Generate four to six AI outputs** for it, using prompts that pull in different directions. Change the mood, the audience, the setting, or the reference point. If you like, add an image or two from far outside your field.
2. **Put them side by side,** all at once, where you can see them together.
3. **Don't pick a winner.** For each one, jot down the thought it sparks: a question, an association, an objection, a feeling.
4. **Look past the object to the layer.** For each output, note what you're actually responding to. Is it the tone? The texture? The structure? The pacing? (One designer Andrés Lucero studied included a photo of a ripe melon on a mood board, not because the project had anything to do with melons, but "to introduce the sense of smell.")
5. **Read the set as a whole.** Write a sentence or two about what the collection suggests that no single output does on its own.

The point isn't to find the best option. It's to notice what the differences between options can spark, and to do the connecting yourself.

8. Who Can Work This Way?

Does working this way take more time? Gloria Chua isn't sure it does.

For her research synthesis, AI "speeds up the first cut. That's the major thing. The most obvious themes become very easy to do." But:

> *"I don't know if it's actually more efficient, because I feel like the interesting thing only happens in the later part of the synthesis, in the second round. Let's say you spend ten hours doing synthesis, before AI and after AI. I feel like it probably took the same amount of time to get to that point. It's just that the interim process looked a bit different."*

Before AI, her team would "get into a room, actually move all the post-its around... draw many diagrams." Now she records her interviews, organizes the points on a digital board, and feeds it all to Claude for a report. "Usually about thirty percent of it is helpful."

She edits the findings she disagrees with, works on it because "it doesn't feel like it's hit the essence," and then spends "the last two to three hours" on her own synthesis, "after various conversations with different team members."

The total time didn't change much. What changed was where her attention went, and she chose how to spend it.

Not everyone feels they have that choice. When I talked with Rina Takikawa, she started to describe how good ideas often come from zoning out and letting them arrive, then caught herself:

> *"But I think that's if you have the privilege of time. And it depends on what the project is. If it's a side project, then sure."*

I pushed back a little. Even if your goal is to be productive, I said, sometimes putting something down and letting it marinate is the most productive thing you can do.

"Yes, I so agree," she said. "But society would tell us otherwise... In a company situation, you can't really say, 'My standup was this, so now I'm just waiting for the creative ideas to come to me.' You have to do something."

I told her about a startup boss I once had, the kind of founder who couldn't stop thinking about the business. Even he would tell people, "You need to put this down and come back to it tomorrow. Do something else." So yes, I said, there's some privilege and circumstance involved. But part of it is just a human management problem.

Takikawa's answer was more personal:

> *"I think it doesn't help if you're the overachiever, workaholic persona that I have. Then it doesn't become a human management problem. It becomes your own management problem, which is something I deal with a lot."*

…

I think part of what makes this hard is that time to think has become scarce in a particular way. Jon McCormack described it as:

> *"not having the headspace to think through things more carefully, or the mental or physical space to do that. Which once was not considered a luxury; it was just considered normal. And now it's probably considered a luxury."*

Rina Takikawa feels that pressure from the inside. As she described in the section on choices, when everyone around you is shipping fast, the pace becomes the new baseline. She's torn about it herself: "As someone so ingrained in startup culture, speed is literally everything in startups. So I go back and forth on this topic as well."

Justin Keoninh made a different choice. He went the traditional route toward design director and then founder, and realized he didn't like it. "What I really enjoyed was getting very curious about a certain problem and spending time with it." He went back to being an individual contributor. "I think that's what a good designer is: really knowing how to sit with a problem and understand its nuance. And I just don't think industries allow that, because of revenue and the incentives of profitability."

Even expertise doesn't fully protect you. Joba Adisa has found that while people with more expertise generally use AI more wisely:

> *"There might be many other things. You might be an expert, and then suddenly you're racing against time... now there's a time crunch, and you don't have the time to sit down and cherry-pick."*

Time isn't the only thing that's unevenly distributed. Another thing is what people already bring.

Fernanda del Monte sees this in the designers who come into her postgraduate program in Mexico City. A few use AI in remarkably creative ways. Most stay in the middle. Many, she said, "don't read, they don't see a lot of cinema, they don't know about semiotics... They don't have abstract thinking."

That matters with AI:

> *"If you cannot think in abstraction, then how can you create a grammar? How can you create a prompt that is the synthesis of an idea?"*

Pann Lim worries about this across a whole generation. With decades of experience, he trusts his own eye. "For those of us who are experienced, and know how to use AI, I think it's fine." But "can you imagine a whole generation that was born with AI already, and hasn't done any critical thinking, any conceptual thinking, in their day-to-day process?"

Punya Mishra makes a related point about writing: beginning writers "haven't found their voice yet," which makes it harder to notice when AI is pulling them away from it.

Research points the same way. In a 2025 study of how employees use AI at work, organizational behavior researchers Shuhua Sun and colleagues found that the creative benefits depended largely on metacognitive strategies: whether people actively thought about the task, noticed the gaps in their own knowledge, and adjusted how they were asking. Just having access to AI wasn't enough.

Looking ahead, the authors raise a worry. They point to what researchers call the Matthew effect, the tendency for advantages to build on themselves, as in "the rich get richer." People who already have strong metacognitive habits "may experience compounding benefits over time," while those without them fall further behind.

But they also note that it could cut the other way, even for people who start out ahead: someone who gets a creative boost from AI "may become increasingly dependent on these tools, potentially reducing autonomy, learning, and networking—factors essential for sustaining creativity over time."

Which way it goes likely depends on how people use the tools, and that's still very much an open question. But the authors are clear on something hopeful: metacognitive strategies "are teachable."

Then there's what gets rewarded.

Even when people understand the risks, the incentives can push the other way. Ken Montgomery's high school students tell him they "understand using LLMs can create cognitive deskilling." But "they also feel they have a lot on their plate. They're getting a grade. They're in a competition: in high school, against other students for college admissions; in college, maybe for job offers."

Pann Lim sees similar pressures in his industry, where expectations have drastically changed because of AI: "A client comes and says, I need 300 social media assets in a month... There's no way you can do it other than using the platforms." Clients now expect "20 versions, 50 versions," even though "sometimes prompting takes even more time... because it gets silly at some point."

When Lim spoke to a group of design students who were, as their lecturers put it, "harshly anti-AI," he didn't tell them they were wrong:

> "No *one can force you to use it, but eventually, if you're looking for a job, and the job requires you to have some form of skill set for it, then it's up to you to decide.*"

And when a field rewards polish alone, it can push originality aside. Justin Keoninh saw it happen in b-boying once the money came in.

"Now the rubric of respect is almost how polished you are, what moves you pull off and execute flawlessly, versus actually extending out from an originality standpoint."

Still, working this way doesn't always cost more time.

Pann Lim's last 10% takes him five minutes by hand, versus an hour of re-prompting. Tyler James sometimes closes the AI tool entirely and doesn't open it again "until there's an actual plan." And Pramod Khambete, after a long session trying to get Claude to help him collect an online archive, finally told it: "if I wanted to do only this, I could have done it myself, just by downloading these pages, in less than half an hour." Claude's reply: "yes, that's perfectly correct."

Sometimes the slower-seeming path is the faster one.

…

I want to be careful here. Some people really do have more time and room to work this way than others. And I'll admit that I sometimes catch myself saying "I don't have time for that" about the very practices I'm describing.

But I consistently find that when I make even a little room, I end up with work that feels more like mine, ideas that go somewhere others' don't, and something closer to Mihaly Csikszentmihalyi's flow (the state where challenge meets skill).

Most people can practice some of this, even if not all of it, all the time.

It's also true that AI can open doors. Pramod Khambete works with NGOs that "can't afford fancy tools." For them, AI does "good-enough stuff… much better than what they could do, including things they could never have imagined they'd be able to do, just through this chat interface." Danny Geisler's father, a longtime game developer, now builds apps by voice, about a thousand of them so far. "He's kind of gotten his creativity back."

Some people, for good reasons, choose not to use AI creatively at all. Grace Liang, a Stanford student and writer, is direct about it: "I never use AI creatively or cognitively." She also points to an imbalance in who gets heard. Online, she sees many creatives arguing against AI in the arts. "However, every time I see big conferences, every time I see big organizations being propped up, it's 'How do we integrate AI into these works?'" And speaking up has its own risks for someone in her position: "Would this lead to more suspicion of me using AI in my own work? Would it be an ideological mark?"

Then there's the question of whose stories get told, and who the tools are built for.

Tyler James warns:

> "*If we aren't intentional about making that access available to everyone, we'd end up in the same place we are today: a very monolithic group telling the same stories, just with different technology.*"

Majo Greloni, who studies AI and the Global South, notes that "all these technologies... were created in the North, and in the South it's always an off-the-shelf thing." Later in our conversation, she added: "The LLM is just based in the States. It probably got most of its information from there. We're probably not going to create solutions for Africa or the Global South."

And Joba Adisa points out that while a student in the US may have a school-issued laptop from kindergarten, the average K-12 student in Nigeria likely doesn't.

...

When Rina Takikawa talked about what she can control, even without the privilege of time, she pointed to her environment: "Who are you surrounded by? What ideas do you consume?" The first question is where we go next. Working with AI can feel like a solo activity. But a lot of what makes a difference in life happens with other people.

Something to try: Go upstream once

Rina Takikawa's other question, "What ideas do you consume?", is one almost anyone can change up, even with very little time. One way to do that is to go upstream of summarized or aggregated information sources you consume.

1. **For a few days, notice how much of what you take in is secondhand**: newsletters, recaps, social media threads, reviews, AI-generated summaries and answers.
2. **Pick one, and go upstream to what it was made from**: the original paper, the full interview, the talk, the book chapter, or the source an AI answer drew on.

3. **Spend some time with the original source.** You don’t have to read every word. Notice what’s there that the summary left out: a caveat, an example, a tone, a detail that changes the meaning.
4. **Afterward, weigh it.** What did going upstream cost you in time and attention? What did you get from the source that the summary didn’t give you? Did the trade feel worthwhile this time? Why or why not?

Some sources might be behind paywalls, so pick one you can actually reach.

You don't have to give up summaries. The goal is to find out, for yourself, what the trade feels like in practice, and to remember that an AI answer is itself a kind of summary of other people's work.

9. Not a Solo Sport

Remember Majo Greloni's teammate, the one who wanted to skip the post-its because AI could give them fifty ideas? There's more to that story.

The competition brought together a team from China, Pakistan, Canada, Germany, and Argentina.

One of her teammates, who might have been very, very, very focused on winning, had loaded "all the context, everything, the papers" into an AI model before they even started. "We're going to make it amazing," he told her. From then on, every new instruction went straight into the model, and the team read what came back.

Greloni felt herself disconnecting.

> "*If we can't have a conversation because the only voice you're going to hear is the LLM, why be here? Why don't we just do individual groups? Why are we six people at this table if we're only going to hear the LLM? Then there was this joke: okay, let's make the LLM the CEO.*"

What bothered her most wasn't the quality of the work.

> "*What a unique opportunity lost, right? Because in the end, who cares if what we're creating is super good? It's also an opportunity to hear each other, understand ways of thinking.*"

And something else was left behind, too.

> "*The methodology, the design that has been worked on for many years, is also left aside... we're just reading the LLMs, which don't know what happens in a brainstorm.*"

Greloni knows what she was missing. Describing a creative retreat she'd been part of, she said:

> "*You feel more alive when you're doing it with other people, because their energy affects you. Things come that you never expected, good or bad, but you have to deal with that. Emotions appear. With the LLMs, not at all.*"

I'd add one more thing to what Greloni said: by running everything first through an LLM, they probably made a winning idea less likely. (Most likely led to less likely in all the wrong ways.)

…

Other people can also help in a way that's easy to overlook: by letting you see their work before it's done.

In the introduction, I said I glaze over when people say you need "taste" without saying how to develop it. James Yu offered one of the more concrete answers I heard on how to develop your own taste, and working with other people is his key.

When he was starting out as a fiction writer, "I spent as much time reading and critiquing other people's work as working on my own." In writing groups, he was "exposed to unfinished works." That mattered, because "when you go to the library or the bookstore, you only see finished works. You don't see the trail, the paper trail of how they got to the final product."

Learning to critique well took him about a year. And it changed how he understood taste:

> "*I think that's what taste is: being able to know those sets of choices, not to get to a story I think is better, but to get to the story I think this author wants to tell. That gap. There's so much of a gap in an unfinished work. I think seeing that gap is taste, and knowing how to traverse that gap as well.*"

His answer reminded me of a warm-up I'd run with university instructors the day before. I asked them about a time they saw a student's work and knew it was unmistakably that student's. One of the most common answers: it didn't feel fully finished. There was something incomplete about it, in a good way.

It also reminded me of how much I loved listening to rehearsal recordings when I played in bands. All those unfinished thoughts made me think a little more… and rock out in the car. (Now that I look back at it, there was something about rocking out to unfinished thoughts on the 280!)

When I told Yu I might start telling my students to read unfinished work, I realized it's a lot of what we already do in design. "Right," he said. "The crits and everything." (Crits are the critique sessions where designers share work in progress.)

…

Why does working with other people matter so much when AI is in the mix?

Part of the answer goes back to the gravitational pull of the most likely. A model blends what it knows into one average. People don't work that way.

As Yuting Deng and her colleagues put it, any one person's way of thinking "can become entrenched and constrain individual ideation through fixation." But because knowledge is spread across many different minds, "the idiosyncratic expertise and perspectives of each individual enable diverse exploration at the collective level."

In other words, each of us gets stuck in our own ways, but we get stuck in different places. That difference is valuable.

It's also exactly what disappears when everyone runs their ideas through the same model, as we saw with Doshi and Hauser's short stories: better individually, more alike collectively. Harsh Kumar's research team is trying to address this: "How can we elicit the uniqueness that everyone has and mix it with LLMs?"

Justin Keoninh sees the value of different perspectives when he hires. The designers who stand out "look at problems in a different way," and when he asks how they got there, "essentially their backgrounds dictated how they looked at problems."

The organizational researcher Sarah Harvey argues that breakthroughs in groups come less from sheer numbers of ideas than from what she calls creative synthesis: "an integration of group members' perspectives into a shared understanding that is unique to the collective."

So where does AI fit? Abel Valdez has a clear view. Instead of each designer having a private, one-on-one conversation with AI, he wants to "put the AI outside," so it can help people make sense of their conversations with each other.

AI becomes something that serves the group's thinking, rather than replacing it. And, he adds, "if the AI is a black box, we should not let the design process be a black box too."

Perspectives can be shared in conversation. The skill of judgment can be harder to pass along.

When I asked people how they knew when something was working, many of them struggled to say. Rina Takikawa told me, "I honestly don't know if I have any tangible advice on how to build the intuition for slowing down."

Jon McCormack, describing how he judges whether a generative system is any good, said, "I've never really tried to formally define what that means. It's just kind of an implicit thing that you do."

Pann Lim put it this way: "I don't know how to explain this human touch. When I see one, I know one."

If the skill of judgment is hard to put into words, how does anyone learn it? Mostly by watching other people use it.

Lim gets this from sitting on award juries:

> "*What I enjoy most about being in the judging room is absorbing all their comments on the work, because that's how learning starts. You hear why they like a piece of work and why they don't, and what the contention was between two different people with creative thinking.*"

Some people are building this kind of openness into how they share AI work. At Tyler James's studio, creators walk others through their projects, sharing the skills they built, their prompts, and their shot guides, "or even things that didn't work, for certain scenes, to really give people a sense of their thought process."

Digital tools can help, because they keep a record. As Sonja Rozental pointed out, "it's not only the lines that are saved, it's the very process of creation that's preserved." Daniel Yin and Danny Geisler's video game builder is built on that idea. Players can "go into the history, find a part of the game you want to start from, fork it, and start building your own game."

Keoninh learns this way from artists he admires. Leonardo is known for the Mona Lisa, “But when you get into his sketches and all these different dimensions, that body of work is very fascinating.” With Virgil Abloh, seeing “his experimentations, his actual process... forms a whole different way of seeing how this person sees the world.” The finished piece shows you what someone made. The process shows you how they decided.

...

Some of what people give each other can’t come from AI at all. Mishra has a name for what’s missing when AI tries to give these kinds of things: a “hollow signal.” A chatbot can say it’s sorry, or tell you your work is great, but there’s nothing behind it. Take patience as an example:

> “*[AI is] infinitely patient, but infinite patience equals infinite indifference, because there’s no cost to its patience. A child is very aware of the cost to the teacher of sitting with them hour after hour.*”

When he described this, I compared a hollow signal to a story without stakes, which isn’t really a story at all. “Exactly,” he said.

Mishra remembers two of his professors from graduate school.

One praised everything he did, even work he knew wasn’t good. “And my thought was, does he think I’m stupid?”

The other, Bill Brewer, told him, “Punya, I’m really disappointed. I expected much better work from you.” Mishra reflected, “Boy, did that make me feel good: that he expects a lot from me.” Brewer had actually seen him. And he had taken a risk to say so, since the relationship could have been damaged. “I recognized the risk he took, and I valued that. We don’t do that for AI systems.”

I don’t want to suggest that more people always means better thinking. Groups can converge on their own in less creative ways.

In a 2011 brainstorming study by cognitive psychologists Nicholas Kohn and Steven Smith of Texas A&M and UT Arlington, people who saw each other’s ideas began matching them. They kept producing just as many ideas, but across a narrower range.

So the question isn't whether to work with others, but which parts of the work go better together and which go better alone.

University of Michigan computational social scientists Joshua Ashkinaze and his colleagues offer one answer. In a 2025 study, participants brainstormed in a chain, each seeing ideas from the people who came before them. In chains of only people, ideas grew more alike over time. When AI-generated ideas were mixed in, the group's ideas stayed more varied. That might seem to contradict what Doshi and Hauser found, but the comparison is different. Compared with people working independently, AI narrowed the range. Compared with a group already drifting toward the same ideas, it widened it.

That suggests an order. Start with your own ideas, alone. Then work together. Bring AI in as one more outside voice when the group starts to sound the same. At Majo Greloni's workshop table, it went the other way: AI showed up before anyone had written down their own ideas, and the conversation never got started.

But even people who value collaboration feel the pull of working alone. László Gaál has watched some AI commercial teams hold on to people like their costume designers and storyboard artists, replacing only the shoot itself. As he put it:

> "*If it's just one person typing the prompts, he won't get the same quality of output as with all these inputs from people who know what they're doing.*"

He wants to make shorts with friends who are cinematographers for the same reason. And yet, he told me, "solo creation is something I really enjoy, having the control."

It's not about working together all the time. But perhaps the default should be to at least ask whether something would be better done together.

Working together with AI in the mix can also feel strange at first, especially if you're used to using it alone. It takes practice.

Something to try: A writers' room with AI in it

Television writers work in writers' rooms, where a group shapes a story by pitching, arguing, and building on each other's ideas. This activity takes inspiration from writers' rooms to explore what it's like to work together with AI in the mix.

1. **Get together with a friend or colleague, and pick something you both care about and are both involved in**: a course you're designing together, a decision your team has to make, a shared project, or a plan for something at home or in your community. It should be a real situation you share, not a hypothetical, and not one person's problem with the other helping out.
2. **Start by writing alone.** Each of you spends five minutes writing down your own take. Don't polish it.
3. **Then come together.** Trade your unfinished takes. Before you say what you'd change, ask what the other person is trying to do. Look for the gap between where their idea is and where they want it to go.
4. **Bring in AI only where you need it.** When you notice the two of you starting to agree too easily, ask a model what neither of you has considered, or for the strongest case against where you've landed. Then set it aside and keep talking.
5. **Then notice and reflect.** Each of you writes one sentence about something the other person saw that the AI didn't, and one sentence about something the AI raised that neither of you had.

Conclusion: Remembering to Be a Designer

I keep thinking about my colleague who forgot to be a designer. What brought them back wasn't a new tool or a better prompt. It was a question: what different ways have you thought of to frame the problem?

Most of this piece has been about questions like that one.

What are you trying to do? What are you noticing, in yourself and in the material? What did you choose, and what got chosen for you?

None of them need to be asked in a particular order. You can start with intent, notice something that changes it, and reflect your way to a new intent. What matters is that you keep asking, and keep experimenting with the answers.

Thinking more doesn't always mean thinking harder. Sometimes it looks a lot more like play. And play can lead to "Yay!" That's not to say thinking more is easy. It's not. But if we tweak our mindset, thinking more can be more fun and fulfilling.

Punya Mishra told me about a game he used to play with his son when he was two. Mishra would say "Humpty," and his son would say "Dumpty." Mishra would say "sat on a," and his son would say "wall." One day, his son twisted himself upside down in his father's arms, and they played again. "Humpty." "Dumpty." "Sat on a..."

> *"He looks at me and says, 'Dumpoony,' and laughs uproariously, as if he's done the funniest thing on earth. That story stayed with me for a long time, and then I realized he was doing autocomplete. But no large language model would ever say 'Dumpoony.'"*

His son knew the right answer. He chose a different one.

> *"He knows that I know he knows the right answer, but he's giving me the wrong answer because he thinks it's funny, and he's looking at me: does Dad think it's funny as well? There's so much happening in that micro-interaction... about agency and intentionality and play."*

Acknowledgments

Thank you to everyone who took the time to talk with me for this project. Every conversation shaped my thinking: Joba Adisa, Giselle Beiguelman, Gloria Chua, Fernanda del Monte, László Gaál, Danny Geisler, Majo Greloni, Tyler James, Justin Keoninh, Pramod Khambete, Harsh Kumar, Grace Liang, Pann Lim, Jon McCormack, Punya Mishra, Ken Montgomery, Sonja Rozental, Advait Sarkar, Rina Takikawa, Samuel Tschepe, Abel Valdez, Vera van der Burg, Daniel Yin, and James Yu.

Thank you to Isabelle Hau, Jessica Tsang, and the entire Stanford Accelerator for Learning team, for the Thought Leadership grant, for the opportunities to share this work, and for the program that nudged me to create it in the first place.

Thank you to Sarah Stein Greenberg and Scott Doorley at the Stanford d.school for their encouragement and support.

And thank you to John Mitchell, whose ongoing collaboration on teaching and learning sparked a lot of the thoughts in this piece, many of which I'm sharing here for the first time.

Bibliography

Alquist, J. L., & Baumeister, R. F. (2024). Learning to love uncertainty. *Current Directions in Psychological Science*, 33(6), 355–360. https://doi.org/10.1177/09637214241279539

Alter, A. L., Oppenheimer, D. M., Epley, N., & Eyre, R. N. (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General*, 136(4), 569–576. https://doi.org/10.1037/0096-3445.136.4.569

Anderson, B. R., Shah, J. H., & Kreminski, M. (2024). Homogenization effects of large language models on human creative ideation. In *Proceedings of the 16th Conference on Creativity & Cognition* (pp. 413–425). Association for Computing Machinery. https://doi.org/10.1145/3635636.3656204

Arnold, K. C., Chauncey, K., & Gajos, K. Z. (2020). Predictive text encourages predictable writing. In *Proceedings of the 25th International Conference on Intelligent User Interfaces* (pp. 128–138). Association for Computing Machinery. https://doi.org/10.1145/3377325.3377523

Ashkinaze, J., Mendelsohn, J., Qiwei, L., Budak, C., & Gilbert, E. (2025). How AI ideas affect the creativity, diversity, and evolution of human ideas: Evidence from a large, dynamic experiment. In *Proceedings of the ACM Collective Intelligence Conference* (pp. 198–213). Association for Computing Machinery. https://doi.org/10.1145/3715928.3737481

Bainbridge, L. (1983). Ironies of automation. *Automatica*, 19(6), 775–779. https://doi.org/10.1016/0005-1098(83)90046-8

Baird, B., Smallwood, J., Mrazek, M. D., Kam, J. W. Y., Franklin, M. S., & Schooler, J. W. (2012). Inspired by distraction: Mind wandering facilitates creative incubation. *Psychological Science*, 23(10), 1117–1122. https://doi.org/10.1177/0956797612446024

Barr, N., Pennycook, G., Stolz, J. A., & Fugelsang, J. A. (2015). The brain in your pocket: Evidence that smartphones are used to supplant thinking. *Computers in Human Behavior*, 48, 473–480. https://doi.org/10.1016/j.chb.2015.02.029

Barto, A., Mirolli, M., & Baldassarre, G. (2013). Novelty or surprise? *Frontiers in Psychology*, 4, Article 907. https://doi.org/10.3389/fpsyg.2013.00907

Basu, A., Duvall, J., & Kaplan, R. (2019). Attention restoration theory: Exploring the role of soft fascination and mental bandwidth. *Environment and Behavior*, 51(9–10), 1055–1081. https://doi.org/10.1177/0013916518774400

Beckman, S. L., & Barry, M. (2007). Innovation as a learning process: Embedding design thinking. *California Management Review*, 50(1), 25–56. https://doi.org/10.2307/41166415

Bellemare-Pepin, A., Lespinasse, F., Thölke, P., Harel, Y., Mathewson, K., Olson, J. A., Bengio, Y., & Jerbi, K. (2026). Divergent creativity in humans and large language models. *Scientific Reports*, 16, Article 1279. https://doi.org/10.1038/s41598-025-25157-3

Benjamin, J. J., Berger, A., Merrill, N., & Pierce, J. (2021). Machine learning uncertainty as a design material: A post-phenomenological inquiry. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (pp. 1–14). Association for Computing Machinery. https://doi.org/10.1145/3411764.3445481

Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.

Bousquette, I. (2026, April 14). Why some companies say AI 'tokenmaxxing' is key to survival. *The Wall Street Journal*. https://www.wsj.com/cio-journal/why-some-companies-say-ai-tokenmaxxing-is-key-to-survival-e699a128

Briñol, P., Petty, R. E., & Tormala, Z. L. (2006). The malleable meaning of subjective ease. *Psychological Science*, 17(3), 200–206. https://doi.org/10.1111/j.1467-9280.2006.01686.x

Buçinca, Z., Malaya, M. B., & Gajos, K. Z. (2021). To trust or to think: Cognitive forcing functions can reduce overreliance on AI in AI-assisted decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), 1–21. https://doi.org/10.1145/3449287

Camarda, A., De Neys, W., Ozkalp-Poincloux, B., Hooge, S., Le Masson, P., Weil, B., & Cassotti, M. (2026). Detecting fixation bias in creative idea generation: Evidence from design novices and experts. *Creativity Research Journal*, 38(1), 189–209. https://doi.org/10.1080/10400419.2024.2424620

Cardoso, C., & Badke-Schaub, P. (2011). The influence of different pictorial representations during idea generation. *The Journal of Creative Behavior*, 45(2), 130–146. https://doi.org/10.1002/j.2162-6057.2011.tb01092.x

Carnat, I. (2024). Human, all too human: Accounting for automation bias in generative large language models. *International Data Privacy Law*. Advance online publication. https://doi.org/10.1093/idpl/ipae018

Carpenter, S. K., Wilford, M. M., Kornell, N., & Mullaney, K. M. (2013). Appearances can be deceiving: Instructor fluency increases perceptions of learning without increasing actual learning. *Psychonomic Bulletin & Review*, 20(6), 1350–1356. https://doi.org/10.3758/s13423-013-0442-z

Chan, J., Dow, S. P., & Schunn, C. D. (2015). Do the best design ideas (really) come from conceptually distant sources of inspiration? *Design Studies*, 36, 31–58. https://doi.org/10.1016/j.destud.2014.08.001

Chandra, K., Li, T.-M., Tenenbaum, J. B., & Ragan-Kelley, J. (2024). Storytelling as inverse inverse planning. *Topics in Cognitive Science*, 16(1), 54–70. https://doi.org/10.1111/tops.12710

Chase, C. C., Malkiewich, L., & Kumar, A. S. (2019). Learning to notice science concepts in engineering activities and transfer situations. *Science Education*, 103(2), 440–471. https://doi.org/10.1002/sce.21496

Chayka, K. (2025, June 25). A.I. is homogenizing our thoughts. *The New Yorker*. https://www.newyorker.com/culture/infinite-scroll/ai-is-homogenizing-our-thoughts

Chen, H., & Jiao, K. (2026). *Central tendency bias in human selection of AI-generated design variations* (arXiv:2607.09018). arXiv. https://doi.org/10.48550/arXiv.2607.09018

Chiang, T. (2024, August 31). Why A.I. isn't going to make art. *The New Yorker*. https://www.newyorker.com/culture/the-weekend-essay/why-ai-isnt-going-to-make-art

Chirayath, G., Premamalini, K., & Joseph, J. (2025). Cognitive offloading or cognitive overload? How AI alters the mental architecture of coping. *Frontiers in Psychology*, 16, Article 1699320. https://doi.org/10.3389/fpsyg.2025.1699320

Christensen, B. T., & Schunn, C. D. (2007). The relationship of analogical distance to analogical function and preinventive structure: The case of engineering design. *Memory & Cognition*, 35(1), 29–38. https://doi.org/10.3758/bf03195939

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Collins, K. M., Sucholutsky, I., Bhatt, U., Chandra, K., Wong, L., Lee, M., Zhang, C. E., Zhi-Xuan, T., Ho, M., Mansinghka, V., Weller, A., Tenenbaum, J. B., & Griffiths, T. L. (2024). Building machines that learn and think with people. *Nature Human Behaviour*, 8(10), 1851–1863. https://doi.org/10.1038/s41562-024-01991-9

Conaway, E. (2026, June 2). *Finding balance in the era of tokenmaxxing*. DORA. https://dora.dev/insights/finding-balance-in-the-era-of-tokenmaxxing/

Cromwell, J. R. (2024). How combinations of constraint affect creativity: A new typology of creative problem solving in organizations. *Organizational Psychology Review*, 14(1), 3–24. https://doi.org/10.1177/20413866231202031

Csikszentmihalyi, M. (1990). *Flow: The psychology of optimal experience*. Harper & Row.

Csikszentmihalyi, M., & Getzels, J. W. (1971). Discovery-oriented behavior and the originality of creative products: A study with artists. *Journal of Personality and Social Psychology*, 19(1), 47–52. https://doi.org/10.1037/h0031106

Danry, V., Pataranutaporn, P., Mao, Y., & Maes, P. (2023). Don't just tell me, ask me: AI systems that intelligently frame explanations as questions improve human logical discernment accuracy over causal AI explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). Association for Computing Machinery. https://doi.org/10.1145/3544548.3580672

David, L., Vassena, E., & Bijleveld, E. (2024). The unpleasantness of thinking: A meta-analytic review of the association between mental effort and negative affect. *Psychological Bulletin*. Advance online publication. https://doi.org/10.1037/bul0000443

Davis, M. S. (1971). That's interesting! Towards a phenomenology of sociology and a sociology of phenomenology. *Philosophy of the Social Sciences*, 1(2), 309–344. https://doi.org/10.1177/004839317100100211

Dekker, S. W. A., & Woods, D. D. (2002). MABA-MABA or abracadabra? Progress on human-automation co-ordination. *Cognition, Technology & Work*, 4(4), 240–244. https://doi.org/10.1007/s101110200022

Deng, Y., Brucks, M., & Toubia, O. (2026). *Examining and addressing barriers to diversity in LLM-generated ideas* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.6332039

Diemand-Yauman, C., Oppenheimer, D. M., & Vaughan, E. B. (2011). Fortune favors the bold (and the italicized): Effects of disfluency on educational outcomes. *Cognition*, *118*(1), 111–115. https://doi.org/10.1016/j.cognition.2010.09.012

DiStefano, P. V., & Beaty, R. (2026). Beyond idea generation: The importance of idea evaluation in human-AI collaborative creativity. In M. J. Worwood & J. C. Kaufman (Eds.), *Generative artificial intelligence and creativity: Precautions, perspectives, and possibilities* (pp. 27–36). Academic Press. https://doi.org/10.1016/b978-0-443-34073-4.00013-7

Dorst, K. (2015). *Frame innovation: Create new thinking by design*. MIT Press.

Dorst, K., & Cross, N. (2001). Creativity in the design process: Co-evolution of problem–solution. *Design Studies*, *22*(5), 425–437. https://doi.org/10.1016/s0142-694x(01)00009-6

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, *10*(28), Article eadn5290. https://doi.org/10.1126/sciadv.adn5290

Dow, S. P., Glassco, A., Kass, J., Schwarz, M., Schwartz, D. L., & Klemmer, S. R. (2010). Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Transactions on Computer-Human Interaction*, *17*(4), 1–24. https://doi.org/10.1145/1879831.1879836

Engeler, N. C., & Gilbert, S. J. (2020). The effect of metacognitive training on confidence and strategic reminder setting. *PLOS ONE*, *15*(10), Article e0240858. https://doi.org/10.1371/journal.pone.0240858

Fisher, M., Goddu, M. K., & Keil, F. C. (2015). Searching for explanations: How the internet inflates estimates of internal knowledge. *Journal of Experimental Psychology: General*, *144*(3), 674–687. https://doi.org/10.1037/xge0000070

Fisher, M., & Oppenheimer, D. M. (2021). Harder than you think: How outside assistance leads to overconfidence. *Psychological Science*, *32*(4), 598–610. https://doi.org/10.1177/0956797620975779

Gabora, L., & Kaufman, S. B. (2010). Evolutionary approaches to creativity. In J. C. Kaufman & R. J. Sternberg (Eds.), *The Cambridge handbook of creativity* (pp. 279–300). Cambridge University Press. https://doi.org/10.1017/CBO9780511763205.018

Garner, S., & McDonagh-Philp, D. (2001). Problem interpretation and resolution via visual stimuli: The use of 'mood boards' in design education. *Journal of Art & Design Education*, *20*(1), 57–64. https://doi.org/10.1111/1468-5949.00250

Gerber, E., & Carroll, M. (2012). The psychological experience of prototyping. *Design Studies*, *33*(1), 64–84. https://doi.org/10.1016/j.destud.2011.06.005

Gerlich, M. (2025). From offloading to engagement: An experimental study on structured prompting and critical reasoning with generative AI. *Data*, *10*(11), Article 172. https://doi.org/10.3390/data10110172

Getzels, J. W., & Csikszentmihalyi, M. (1976). *The creative vision: A longitudinal study of problem finding in art*. Wiley.

Ghafouri, B. (2025). *The variance paradox: How AI reduces diversity but increases novelty* (arXiv:2508.19264). arXiv. https://doi.org/10.48550/arXiv.2508.19264 (Earlier version titled *A theory of information, variation, and artificial intelligence*.)

Gilbert, S. J., Bird, A., Carpenter, J. M., Fleming, S. M., Sachdeva, C., & Tsai, P.-C. (2020). Optimal use of reminders: Metacognition, effort, and cognitive offloading. *Journal of Experimental Psychology: General*, *149*(3), 501–517. https://doi.org/10.1037/xge0000652

Glenberg, A. M., Wilkinson, A. C., & Epstein, W. (1982). The illusion of knowing: Failure in the self-assessment of comprehension. *Memory & Cognition*, *10*(6), 597–602. https://doi.org/10.3758/bf03202442

Grinschgl, S., Papenmeier, F., & Meyerhoff, H. S. (2021). Consequences of cognitive offloading: Boosting performance but diminishing memory. *Quarterly Journal of Experimental Psychology*, *74*(9), 1477–1496. https://doi.org/10.1177/17470218211008060

Guo, Y., Lin, S., Acar, S., Jin, S., Xu, X., Feng, Y., & Zeng, Y. (2022). Divergent thinking and evaluative skill: A meta-analysis. *The Journal of Creative Behavior*, *56*(3), 432–448. https://doi.org/10.1002/jocb.539

Hao, Q., Xu, F., Li, Y., & Evans, J. (2026). Artificial intelligence tools expand scientists' impact but contract science's focus. *Nature*, *649*(8099), 1237–1243. https://doi.org/10.1038/s41586-025-09922-y

Harvey, S. (2014). Creative synthesis: Exploring the process of extraordinary group creativity. *Academy of Management Review*, *39*(3), 324–343. https://doi.org/10.5465/amr.2012.0224

Henkel, L. A. (2014). Point-and-shoot memories: The influence of taking photos on memory for a museum tour. *Psychological Science*, *25*(2), 396–402. https://doi.org/10.1177/0956797613504438

Henry, D. (2023, December 9). In 1977, a 12-year-old invented record scratching and changed hip-hop forever [Audio]. NPR. https://www.npr.org/2023/12/09/1218415793/in-1977-a-12-year-old-invented-record-scratching-and-changed-hip-hop-forever

Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, *7*(2), 174–196. https://doi.org/10.1145/353485.353487

Huang, M., Jin, M., & Li, N. (2024). *Augmenting minds or automating skills: The differential role of human capital in generative AI's impact on creative tasks* (arXiv:2412.03963). arXiv. https://doi.org/10.48550/arXiv.2412.03963

Jansson, D. G., & Smith, S. M. (1991). Design fixation. *Design Studies*, *12*(1), 3–11. https://doi.org/10.1016/0142-694x(91)90003-f

Jose, B., Cherian, J., Verghis, A. M., Varghise, S. M., S, M., & Joseph, S. (2025). The cognitive paradox of AI in education: Between enhancement and erosion. *Frontiers in Psychology*, *16*, Article 1550621. https://doi.org/10.3389/fpsyg.2025.1550621

Kapur, M. (2008). Productive failure. *Cognition and Instruction*, *26*(3), 379–424. https://doi.org/10.1080/07370000802212669

Kohn, N. W., & Smith, S. M. (2011). Collaborative fixation: Effects of others' ideas on brainstorming. *Applied Cognitive Psychology*, *25*(3), 359–371. https://doi.org/10.1002/acp.1699

Kolko, J. (2010). Abductive thinking and sensemaking: The drivers of design synthesis. *Design Issues*, *26*(1), 15–28. https://doi.org/10.1162/desi.2010.26.1.15

Kornish, L. J., & Jones, S. M. (2021). Raw ideas in the fuzzy front end: Verbosity increases perceived creativity. *Marketing Science*. Advance online publication. https://doi.org/10.1287/mksc.2021.1300

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:2506.08872). arXiv. https://doi.org/10.48550/arXiv.2506.08872

Kruger, J., Wirtz, D., Van Boven, L., & Altermatt, T. W. (2004). The effort heuristic. *Journal of Experimental Social Psychology*, 40(1), 91–98. https://doi.org/10.1016/s0022-1031(03)00065-9

Kumar, H., Vincentius, J., Jordan, E., & Anderson, A. (2025). Human creativity in the age of LLMs: Randomized experiments on divergent and convergent thinking. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems* (pp. 1–18). Association for Computing Machinery. https://doi.org/10.1145/3706598.3714198

Kusumegi, K., Yang, X., Ginsparg, P., de Vaan, M., Stuart, T., & Yin, Y. (2025). Scientific production in the era of large language models. *Science*, 390(6779), 1240–1243. https://doi.org/10.1126/science.adw3000

Lee, D., Pruitt, J., Zhou, T., Du, J., & Odegaard, B. (2025). Metacognitive sensitivity: The key to calibrating trust and optimal decision making with AI. *PNAS Nexus*, 4(5), Article pgaf133. https://doi.org/10.1093/pnasnexus/pgaf133

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* (pp. 1–23). Association for Computing Machinery. https://doi.org/10.1145/3706598.3713778

Leroy, S. (2009). Why is it so hard to do my work? The challenge of attention residue when switching between work tasks. *Organizational Behavior and Human Decision Processes*, 109(2), 168–181. https://doi.org/10.1016/j.obhdp.2009.04.002

Linden, K., Hachem, H.-H., & Kondyli, V. (2025). Homo promptus: Predicting the impact of generative AI on human memory and creativity. *Memory, Mind & Media*, 4, Article e15. https://doi.org/10.1017/mem.2025.10012

Liu, G., Christian, B., Dumbalska, T., Bakker, M. A., & Dubey, R. (2026). *AI assistance reduces persistence and hurts independent performance* (arXiv:2604.04721). arXiv. https://doi.org/10.48550/arXiv.2604.04721

Lucero, A. (2012). Framing, aligning, paradoxing, abstracting, and directing: How design mood boards work. In *Proceedings of the Designing Interactive Systems Conference* (pp. 438–447). Association for Computing Machinery. https://doi.org/10.1145/2317956.2318021

Madjar, N., & Shalley, C. E. (2008). Multiple tasks' and multiple goals' effect on creativity: Forced incubation or just a distraction? *Journal of Management*, 34(4), 786–805. https://doi.org/10.1177/0149206308318611

Marche, S. (2023, May 9). The future of writing is a lot like hip-hop. *The Atlantic*. https://www.theatlantic.com/technology/archive/2023/05/generative-ai-novel-writing-experiment-stephen-marche/673997/

Martin, L., & Schwartz, D. L. (2014). A pragmatic perspective on visual representation and creative thinking. *Visual Studies*, *29*(1), 80–93. https://doi.org/10.1080/1472586x.2014.862997

Masood, A. (2025, August 4). *The outsourced mind: Navigating the risks and rewards of cognitive offloading*. Medium. https://medium.com/@adnanmasood/the-outsourced-mind-navigating-the-risks-and-rewards-of-cognitive-offloading-9e1e70ee2efb

McCormack, J. (2017). Working with generative systems: An artistic perspective. In *Proceedings of the Conference on Electronic Visualisation and the Arts* (EVA '17) (pp. 213–218). https://doi.org/10.14236/ewic/eva2017.47

Meincke, L., Mollick, E. R., & Terwiesch, C. (2024). *Prompting diverse ideas: Increasing AI idea variance* [Preprint]. SSRN. https://doi.org/10.2139/ssrn.4708466

Meincke, L., Nave, G., & Terwiesch, C. (2025). ChatGPT decreases idea diversity in brainstorming. *Nature Human Behaviour*, *9*(6), 1107–1109. https://doi.org/10.1038/s41562-025-02173-x

Moon, K., Green, A. E., & Kushlev, K. (2025). Homogenizing effect of large language models (LLMs) on creative diversity: An empirical comparison of human and ChatGPT writing. *Computers in Human Behavior: Artificial Humans*, 6, Article 100207. https://doi.org/10.1016/j.chbah.2025.100207

Moon, K., Johnson, D. R., Kushlev, K., Lira, B., Viskontas, I., Kaufman, J. C., Beaty, R. E., Duckworth, A. L., & Green, A. E. (2026). *Thinking outside the bots: Human ideas in AI's semantic blind spots are not inferior* (arXiv:2610.17324). arXiv. https://doi.org/10.48550/arXiv.2610.17324

Moon, K., Kushlev, K., Bank, A., Lira, B., Viskontas, I., Kaufman, J. C., Johnson, D. R., Duckworth, A. L., & Green, A. E. (2025). *The creative link between words and ideas is weakening in the AI era* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/jsz58_v5

Mueller, J. S., Melwani, S., & Goncalo, J. A. (2012). The bias against creativity: Why people desire but reject creative ideas. *Psychological Science*, *23*(1), 13–17. https://doi.org/10.1177/0956797611421018

Oakley, B., Johnston, M., Chen, K.-Z., Jung, E., & Sejnowski, T. (in press). The memory paradox: Why our brains need knowledge in an age of AI. In *The artificial intelligence revolution: Challenges and opportunities*. Springer Nature. Preprint: https://doi.org/10.2139/ssrn.5250447

Ong, W. J. (1986). Writing is a technology that restructures thought. In G. Baumann (Ed.), *The written word: Literacy in transition* (pp. 23–50). Clarendon Press.

Oppenheimer, D. M. (2006). Consequences of erudite vernacular utilized irrespective of necessity: Problems with using long words needlessly. *Applied Cognitive Psychology*, *20*(2), 139–156. https://doi.org/10.1002/acp.1178

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, *12*(6), 237–241. https://doi.org/10.1016/j.tics.2008.02.014

Oppenheimer, D. M., Cash, T. N., & Connell Pensky, A. E. (2025). You've got AI friend in me: LLMs as collaborative learning partners. *International Journal of Artificial Intelligence in Education*, *35*(6), 3896–3921. https://doi.org/10.1007/s40593-025-00521-6

Oppenheimer, D. M., & Patterson, M. T. (2025). Thinking outside the box means thinking outside the search engine. *Memory & Cognition*, *53*(8), 2686–2699. https://doi.org/10.3758/s13421-025-01732-x

Pandey, P., Singh, T., & Kumar, S. (2023). Cognitive offloading: Systematic review of a decade. *International Journal of Indian Psychology*, *11*(2), 1545–1563. https://doi.org/10.25215/1102.163

Papert, S. (1987). Computer criticism vs. technocentric thinking. *Educational Researcher*, *16*(1), 22–30. https://doi.org/10.2307/1174251

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *52*(3), 381–410. https://doi.org/10.1177/0018720810376055

Pea, R. D. (1987). The aims of software criticism: Reply to Professor Papert. *Educational Researcher*, *16*(5), 4–8. https://doi.org/10.3102/0013189x016005004

Perkins, D. N., & Salomon, G. (2012). Knowledge to go: A motivational and dispositional view of transfer. *Educational Psychologist*, *47*(3), 248–258. https://doi.org/10.1080/00461520.2012.693354

Polanyi, M. (1966). *The tacit dimension*. Doubleday.

Poulsgaard, K. S., & Malafouris, L. (2023). Understanding the hermeneutics of digital materiality in contemporary architectural modelling: A material engagement perspective. *AI & Society*, *38*(6), 2217–2227. https://doi.org/10.1007/s00146-020-01044-5

Rietzschel, E. F., Nijstad, B. A., & Stroebe, W. (2010). The selection of creative ideas after individual idea generation: Choosing between creativity and impact. *British Journal of Psychology*, *101*(1), 47–68. https://doi.org/10.1348/000712609x414204

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

Rolfe, J. M. (1982). Determining the training effectiveness of flight simulators: Some basic issues. *The Aeronautical Journal*, *86*(853), 85–86. https://doi.org/10.1017/s0001924000018534

Rooseboom, H. (2006). Myths and misconceptions: Photography and painting in the nineteenth century (J. Rudge, Trans.). *Simiolus: Netherlands Quarterly for the History of Art*, *32*(4), 291–313. https://www.jstor.org/stable/20355339

Rozental, S., & de Rooij, A. (2026). Neither tool nor collaborator: Rethinking human–AI co-creativity in artistic practice with material engagement theory. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems* (pp. 1–20). Association for Computing Machinery. https://doi.org/10.1145/3772318.3791752

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, *1*(1), 7–59. https://doi.org/10.1007/bf00055564

Sarkar, A. (2024). AI should challenge, not obey. *Communications of the ACM*, *67*(10), 18–21. https://doi.org/10.1145/3649404

Schön, D. A. (1983). *The reflective practitioner: How professionals think in action*. Basic Books.

Schwartz, D. L., & Bransford, J. D. (1998). A time for telling. *Cognition and Instruction*, *16*(4), 475–522. https://doi.org/10.1207/s1532690xci1604_4

Sektani, H. H. J., Miran, F. D., Abdullah, H. K., & Sherwani, A. B. (2026). Generative AI and the reordering of design creativity in architectural education: Evidence from undergraduate students. *Architecture*, 6(3), Article 133. https://doi.org/10.3390/architecture6030133

Sengers, P., Boehner, K., David, S., & Kaye, J. (2005). Reflective design. In *Proceedings of the 4th Decennial Conference on Critical Computing: Between Sense and Sensibility* (pp. 49–58). Association for Computing Machinery. https://doi.org/10.1145/1094562.1094569

Shukla, P., Bui, P., Levy, S. S., Kowalski, M., Baigelenov, A., & Parsons, P. (2025). De-skilling, cognitive offloading, and misplaced responsibilities: Potential ironies of AI-assisted design. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems* (pp. 1–7). Association for Computing Machinery. https://doi.org/10.1145/3706599.3719931

Simkute, A., Tankelevitch, L., Kewenig, V., Scott, A. E., Sellen, A., & Rintel, S. (2024). Ironies of generative AI: Understanding and mitigating productivity loss in human-AI interaction. *International Journal of Human–Computer Interaction*. Advance online publication. https://doi.org/10.1080/10447318.2024.2405782

Singh, A., Taneja, K., Guan, Z., & Ghosh, A. (2025). *Protecting human cognition in the age of AI* (arXiv:2502.12447). arXiv. https://doi.org/10.48550/arXiv.2502.12447

Slamecka, N. J., & Graf, P. (1978). The generation effect: Delineation of a phenomenon. *Journal of Experimental Psychology: Human Learning and Memory*, 4(6), 592–604. https://doi.org/10.1037/0278-7393.4.6.592

Soderstrom, N. C., & Bjork, R. A. (2015). Learning versus performance: An integrative review. *Perspectives on Psychological Science*, 10(2), 176–199. https://doi.org/10.1177/1745691615569000

Sourati, Z., Ziabari, A. S., & Dehghani, M. (2026). The homogenizing effect of large language models on human expression and thought. *Trends in Cognitive Sciences*, 30(9), 805–816. https://doi.org/10.1016/j.tics.2026.01.003

Sparrow, B., & Chatman, L. (2013). Social cognition in the internet age: Same as it ever was? *Psychological Inquiry*, 24(4), 273–292. https://doi.org/10.1080/1047840x.2013.827079

Sparrow, B., Liu, J., & Wegner, D. M. (2011). Google effects on memory: Cognitive consequences of having information at our fingertips. *Science*, 333(6043), 776–778. https://doi.org/10.1126/science.1207745

Stadler, M., Bannert, M., & Sailer, M. (2024). Cognitive ease at a cost: LLMs reduce mental effort but compromise depth in student scientific inquiry. *Computers in Human Behavior*, *160*, Article 108386. https://doi.org/10.1016/j.chb.2024.108386

Stokes, P. D. (2025). Paired constraints: A paradigm for solving the creativity problem. *Possibility Studies & Society*, *3*(3), 454–462. https://doi.org/10.1177/27538699241308851

Storm, B. C., & Stone, S. M. (2015). Saving-enhanced memory: The benefits of saving on the learning and remembering of new information. *Psychological Science*, *26*(2), 182–188. https://doi.org/10.1177/0956797614559285

Sun, S., Li, Z. A., Foo, M.-D., Zhou, J., & Lu, J. G. (2025). How and for whom using generative AI affects creativity: A field experiment. *Journal of Applied Psychology*, *110*(12), 1561–1573. https://doi.org/10.1037/apl0001296

Suzuki, Y., & Banaei-Kashani, F. (2025). *Universe of thoughts: A computational framework for creative reasoning in large language models* (arXiv:2511.20471). arXiv. https://doi.org/10.48550/arXiv.2511.20471 (Earlier version titled *Universe of thoughts: Enabling creative reasoning with large language models*.)

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 the CHI Conference on Human Factors in Computing Systems* (pp. 1–24). Association for Computing Machinery. https://doi.org/10.1145/3613904.3642902

Terwiesch, C., Meincke, L., Girotra, K., Mollick, E., Nave, G., & Ulrich, K. (2026). Artificial intelligence and its impact on creativity and diversity: An empirical study of large language model-generated product ideas. *Production and Operations Management*. Advance online publication. https://doi.org/10.1177/10591478261474243

Tromp, C. (2023). Integrated constraints in creativity: Foundations for a unifying model. *Review of General Psychology*, *27*(1), 41–61. https://doi.org/10.1177/10892680211060027

Tromp, C. (2024). Creativity from constraint exploration and exploitation. *Psychological Reports*, *127*(4), 1818–1843. https://doi.org/10.1177/00332941221114421

Tschepe, S. (2026, March 6). Ich bin nicht immun gegen das, worüber ich schreibe [I am not immune to what I write about]. *Das Kontinuum*. https://daskontinuum.substack.com/p/ich-bin-nicht-immun-gegen-das-woruber

Valkenburg, R., & Dorst, K. (1998). The reflective practice of design teams. *Design Studies*, *19*(3), 249–271. https://doi.org/10.1016/s0142-694x(98)00011-8

van der Burg, V. (2026). *The material gaze: Reflective AI as a design practice* [Doctoral dissertation, Delft University of Technology]. https://doi.org/10.4233/uuid:fb78a0d0-0eb9-46e9-b7af-d54bca8796ca

van der Linden, D., Tops, M., & Bakker, A. B. (2021). Go with the flow: A neuroscientific view on being fully engaged. *European Journal of Neuroscience*, 53(4), 947–963. https://doi.org/10.1111/ejn.15014

Wadinambiarachchi, S., Kelly, R. M., Pareek, S., Zhou, Q., & Velloso, E. (2024). The effects of generative AI on design fixation and divergent thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (pp. 1–18). Association for Computing Machinery. https://doi.org/10.1145/3613904.3642919

Wilson, T. D., Reinhard, D. A., Westgate, E. C., Gilbert, D. T., Ellerbeck, N., Hahn, C., Brown, C. L., & Shaked, A. (2014). Just think: The challenges of the disengaged mind. *Science*, *345*(6192), 75–77. https://doi.org/10.1126/science.1250830

Worwood, M. J. (2026). The convenience trap: How AI could hinder creativity in education. In M. J. Worwood & J. C. Kaufman (Eds.), *Generative artificial intelligence and creativity: Precautions, perspectives, and possibilities* (pp. 179–192). Academic Press. https://doi.org/10.1016/b978-0-443-34073-4.00024-1

Wu, S., Liu, Y., Ruan, M., Chen, S., & Xie, X.-Y. (2025). Human-generative AI collaboration enhances task performance but undermines human's intrinsic motivation. *Scientific Reports*, *15*(1), Article 15105. https://doi.org/10.1038/s41598-025-98385-2

Yan, L., Pammer-Schindler, V., Mills, C., Nguyen, A., & Gašević, D. (2025). Beyond efficiency: Empirical insights on generative AI's impact on cognition, metacognition and epistemic agency in learning. *British Journal of Educational Technology*, 56(5), 1675–1685. https://doi.org/10.1111/bjet.70000

Zhai, C., Wibowo, S., & Li, L. D. (2024). The effects of over-reliance on AI dialogue systems on students' cognitive abilities: A systematic review. *Smart Learning Environments*, *11*, Article 28. https://doi.org/10.1186/s40561-024-00316-7

Zhang, C., Kwon, Y. P., Kramer, J., Kim, E., & Agogino, A. M. (2017). Concept clustering in design teams: A comparison of human and machine clustering. *Journal of Mechanical Design*, *139*(11), Article 111414. https://doi.org/10.1115/1.4037478

Zhang, Z. T., von Felten, N., Reicherts, L., Tankelevitch, L., Guan, Z., Rintel, S., Fu, Y., He, J., Holstein, K., Sarkar, A., Ramos, G., Schmitt, A., Singh, A., Li, H., Palani, S., & Dalsgaard, P. (2026). Tools for thought: Understanding, protecting, and augmenting human cognition with generative AI—From vision to implementation. In *Proceedings of the Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems* (pp. 1–7). Association for Computing Machinery. https://doi.org/10.1145/3772363.3778733