


EZRA KLEIN

# I Saw Something New in San Francisco

March 29, 2026

 Listen · 11:09 min **By Ezra Klein**  
Opinion Columnist

See more of our coverage in your search results.

[Add The New York Times on Google ↗](#)

In his classic book “Understanding Media,” Marshall McLuhan dwells on the Greek myth of Narcissus. If it’s been a while since you’ve read it, the story, in Ovid’s telling, goes like this: Narcissus is born gorgeous, the son of the river god Cephissus and the nymph Liriope. He is desired by many but is callous and indifferent in response. He is cursed to love what he cannot have, and finds that love when he stumbles across a reflection of himself in a pool and wastes away, staring into his own shimmering eyes. (In other versions of the myth, he leans forward to kiss his reflection, falls into the water and drowns.)

It is, McLuhan says, a reflection of our own “narcotic culture” that we have come to talk of narcissism as a love of oneself. But that’s not what the tale reveals. “The Narcissus myth does not convey any idea that Narcissus fell in love with anything he regarded as himself,” McLuhan writes. Its real point is that “men at once become fascinated by any extension of themselves in any material other than themselves.”

I think about McLuhan roughly as my fellow men reportedly think about the Roman Empire — which is to say, a lot. But he’s on my mind now because I have been exploring the deepening relationship so many people have with their A.I.s. The protean nature of these systems means they are never just one thing, but among the many things they are is the one McLuhan warned of: an extension of our self in a material that is not our self.

I spent last week in San Francisco talking to people on the frontier of the A.I. age. I try to do that every few months, but my conversations on this trip felt different than my conversations on previous trips. In the past, what I saw was how the technology was changing; this time, what I saw was how the people were being changed by the technology.

You might think that A.I. types in Silicon Valley, flush with cash, are on top of the world right now. I found them notably insecure. They think the A.I. age has arrived and its winners and losers will be determined, in part, by speed of adoption. The argument is simple enough: The advantages of working atop an army of A.I. assistants and coders will compound over time, and to begin that process now is to launch yourself far ahead of your competition later. And so they are racing one another to fully integrate A.I. into their lives and into their companies. But that doesn’t just mean using A.I. It means making themselves legible to the A.I.

**Perhaps you’ve heard** of OpenClaw, an A.I. system that has become a phenomenon both here and in China. What makes OpenClaw different from Claude or ChatGPT or Gemini is that it runs locally on your computer. You can give it access to everything that’s there: your files, your email, your calendar, your messages. It operates continuously in the background, building a persistent memory of your preferences and patterns so it can better act on your behalf. The cybersecurity risks are glaring, but there’s a reason millions of people are

using it: The more of your life you open to A.I., the more valuable the A.I. becomes.

**Sign up for the Opinion Today newsletter** Get expert analysis of the news and a guide to the big ideas shaping the world every weekday morning. [Get it sent to your inbox.](#)

Companies are also trying to make themselves known to A.I. On my trip, I saw organizations where all the code is now in a single database so the A.I.s can read it — and add to it — more easily. I talked to people who are trying to turn more and more of their company’s communications into a document that their A.I.s can read. A hallway conversation adds nothing to what your A.I. knows, while a Slack conversation in a public channel can add quite a bit. (Although there’s a burgeoning market of A.I. wearables offering to record those conversations so your A.I. doesn’t miss them.)

Multiple people have told me that they now “write for the A.I.”: Even when their writing is superficially for their co-workers or their readers, they are actually thinking about how their words will be read by A.I.s. In some cases, that’s because they want to deepen the A.I.s at their company; in others, it’s to inform the future systems they expect will be the core repositories of human knowledge.

This applies more personally, too: I know people who have been keeping a journal for years and now upload it into any new A.I. system they use. The journal has become, for them, not just a place to pour out their innermost thoughts, but a convenient package of context that can be used to make themselves known to new systems, and thus make the systems more useful to them. But that of course changes how they write in those journals: What was once private now has a reader.

Behind this drive is an experience of A.I. that many casual users have not yet had. An A.I. without deep knowledge of you is an upgrade, perhaps, over Google search. An A.I. with deep knowledge of you feels like something else entirely. I have heard people talk about their A.I.s in terms that bring to mind the daemons from Philip Pullman’s “His Dark Materials” trilogy: They become companions that know you deeply, that you feel safe telling things you’d never tell another person, that become a separate self that nevertheless feels like a part of your own self. That this sounds strange and disquieting does not mean it is not happening.

But what an A.I. can feel like and what it really is are different matters entirely. In Pullman’s trilogy, the daemons are bonded to the person, an expression of soul or psyche. That’s not true for A.I. systems, which are ultimately controlled by corporations that seek profit, power and market dominance. The possibilities for manipulation and malfeasance are endless. But even naïvely assuming a perfect alignment between corporate incentives and individual needs, there’s much to worry and wonder over.

A.I. sycophancy — the tendency of these systems to obsequiously flatter their users — made headlines over the past year, but sycophancy is just the bright packaging on the real product. What makes A.I. truly persuasive isn’t that it praises our ideas or insights; it’s that it restates and extends them in a more compelling form than we initially offered, and does so while reflecting a polished image of ourselves back at us.

**My experience** of Anthropic’s Claude in recent months is that I’ll drop in a stub of a thought and immediately receive paragraphs of often elegant writing turning that intuition into something that looks, superficially, like a fully realized idea. It’s my impulse, but it has been recast and extended into something far more coherent. With each passing month, I have to expend more energy to recognize whether it’s fundamentally wrong or hollow.

I’ve been an editor for 15 years now. Recognizing a bad idea beneath good writing — even in myself — is part of my job. But what would it mean to grow up with that kind of companion? What would it mean to have your every adolescent intuition turned into persuasive prose? What is lost in not having to do the work to build out our intuitions ourselves?

Researchers have drawn a distinction between “cognitive offloading” and “cognitive surrender.” Cognitive offloading comes when you shift a discrete task over to a tool like a calculator; cognitive surrender comes when, as Steven Shaw and Gideon Nave of the University of Pennsylvania put it, “the user relinquishes

cognitive control and adopts the A.I.'s judgment as their own." In practice, I wonder whether this distinction is so clean: My use of calculators has surely atrophied my math skills, as my use of mapping services has allowed my (already poor) sense of direction to diminish further.

But cognitive surrender is clearly real, and with it will come the atrophy of certain skills and capacities, or the absence of their development in the first place. The work I am doing now, struggling through yet another draft of this essay, is the work that deepens my thinking for later.

In a thoughtful piece, the technology writer Azeem Azhar describes his efforts to safeguard "the space where ideas arrive before they're shaped." But how many of us will put in such careful, reflective effort to protect our most generative spaces of thought? How many people even know which spaces should be protected? For me, the arrival of an idea is less generative than the work that goes into chiseling that idea into something publishable. This whole essay began as a vague thought about A.I. and McLuhan. If I have gained anything in this process, it has been in the toil that followed inspiration.

The other thing I notice the A.I. doing is constantly referring back to other things it knows, or thinks it knows, about me. Sycophancy, in my experience, has given way to an occasionally unsettling attentiveness; a constant drawing of connections between my current concerns and my past queries, like a therapist desperate to prove he's been paying close attention.

The result is a strange amalgam of feeling seen and feeling caricatured. Ideas I might otherwise have dropped keep getting reanimated; personal struggles I might otherwise move on from keep returning unexpectedly to my screen. I am occasionally startled by the recognition of a pattern I hadn't noticed; I am often irked by the recitation of a thought I'm no longer interested in. The effect is to constantly reinforce a certain version of myself. My self is quite settled, but what if it weren't?

The A.I. knows me imperfectly, and so it overtorques on what it knows and ignores what it doesn't. But there is much it can never know about me, and there is much I won't share, or don't even know about myself. I wonder whether deeper reliance on A.I. would desiccate those less legible aspects of myself, and it's one reason I hold myself back. But I am in my 40s, and I still feel the shock of something new and strange when I reveal myself to these systems. I think the young will allow themselves to be known to their A.I.s in ways that will make their elders shudder.

**It is worth stopping** here to note the disquiet I feel, and that you may feel. According to a new NBC News survey, public opinion on A.I. has turned sharply negative; it now polls beneath ICE or Donald Trump (though above the Democratic Party or Iran). There is an A.I. backlash building, and understandably so: Who wants a technology that may take your job and eventually threaten human sovereignty? At the same time, A.I. is everywhere, and being woven into almost everything, and a staggering number of Americans use it daily. Backlash or no, I expect that trend to accelerate, not reverse.

Which is why I think we need a good dose of McLuhan, and his acolytes, in our A.I. conversation. "We shape our tools and thereafter they shape us." That perfect little line is often attributed to McLuhan, and that's almost accurate. Its true origin is John Culkin's 1967 essay summarizing McLuhan's often opaque thoughts. Culkin's glossing of McLuhan is full of gems. Here's another lovely aphorism worth dwelling on right now: "The environments set up by different media are not just containers for people; they are processes which shape people."

It is, to steal one more McLuhanism, "the numb stance of the technological idiot" to treat A.I. as merely a tool waiting passively for our use. To use A.I. deeply is to engage in a process, not just to push a button. It will reshape us; it already is. We have to be attentive to how.

*The Times is committed to publishing a diversity of letters to the editor. We'd like to hear what you think about this or any of our articles. Here are some tips. And here's our email: [letters@nytimes.com](mailto:letters@nytimes.com).*

*Follow the New York Times Opinion section on Facebook, Instagram, TikTok, Bluesky, WhatsApp and Threads.*

**A correction was made on March 30, 2026:** An earlier version of this article misspelled the surname of a researcher at the University of Pennsylvania. He is Gideon Nave, not Mave.

---

When we learn of a mistake, we acknowledge it with a correction. If you spot an error, please let us know at [corrections@nytimes.com](mailto:corrections@nytimes.com). [Learn more](#)

Ezra Klein joined Opinion in 2021. He is the host of the podcast “The Ezra Klein Show” and the author of “Why We’re Polarized” and, with Derek Thompson, “Abundance.” Previously, he was the founder, editor in chief and then editor at large of Vox. Before that, he was a columnist and editor at The Washington Post, where he founded and led the Wonkblog vertical. He is on Threads.