

TECHNOLOGY

The Feeling of Control Slipping Away

AI is causing a crisis of agency.

By Charlie Warzel



Illustration by The Atlantic. Source: Solarseven / Getty.

MAY 30, 2026, 7 AM ET

SHARE AS GIFT 

DISCUSS 136

REMOVE 

Back in the web-traffic-obsessed days of 2018, at a time of dawning awareness of how easily audiences online could be manipulated and spoofed by bots, the writer Max Read argued that the internet had crossed a threshold known as “the Inversion.” Not only had bots proliferated across the internet; they had come to constitute it. In outnumbering humans, bots were also loosening everyone’s grasp on the very reality of online experience. “What’s gone from the internet, after all, isn’t ‘truth,’ but trust: the sense that the people and things we encounter are what they represent themselves to be,” Read wrote.

Today, “the Inversion” feels almost quaint. Autonomous AI agents roam the internet,

answering emails, sending texts, and occasionally deleting the code repositories of entire companies. An endless library of chatbot-speak crowds out human-written words in every Google search. Bots are spinning up music and videos, conjuring bad poetry and prose, building websites, doing research, making transactions, writing plodding memos to your boss, solving geometry conjectures. Those AI outputs then ride the rails of an internet controlled by black-box algorithms. Computers talk to computers, producing information to train computers to sound more like humans or to better engage them. Humans type into the box, scroll, and wait.

AI is driving people insane in all kinds of ways. Its overwhelming speed and existential stakes have given rise to generalized malaise and hostility directed at the industry, to say nothing of actual cases of AI psychosis. But a lot of this is subtler—a deepening of the bewildering, corrosive feeling Read previously described. Culturally, the flood of slop, AI influencers, fake accounts, and AI tools is blurring the lines of an already post-truth age. A specific paranoia is in the air, an abiding concern about being manipulated, suckered, influenced. Stealth marketing campaigns, mercenary armies of bots, and paid clippers have led anyone or anything that appears dubious to be deemed a potential “psyop.” Cheap imitations of expressions of human creativity are easier than ever to fake. Sentiment, perhaps even popularity, is easier to manipulate. On top of all this is the push into agentic AI—a future we’re told will consist of an internet crammed with bots performing human tasks.

People who don’t feel empowered by all of this are unmoored. Across so many levels of culture, there’s a feeling of control slipping ever so slightly away. You, me, all of us, whether or not we enjoy or use these tools, are living through a crisis of agency. The agita and paranoia, even the excitement—over AI’s encroachment on work, education, art, and culture—is the by-product of a cultural and technological moment in which humans are sliding into a more passive role in many activities. One way to look at the generative-AI boom is as a massive societal experiment foisted on us by Silicon Valley, the animating question of which is: What is a human for?

When you start looking, you see the anxiety over agency everywhere. You see it in the reactions to the mass layoffs at places such as Meta in preparation for an AI transformation, in the coverage of venture-capital-funded, bulk-content-creation bot-army start-ups that proudly claim, “Never pay a human again.” You can sense it among the software developers who feel that their reliance on coding tools is eroding their skill set, in the executives who confess that they don’t know whether their AI

spend is justifiable. Or when you read reports that medical journals are filling with made-up citations, or a study that suggests that chatbot use is degrading our thinking, or an announcement from Google that it will offer an alternative to its link-based search results: AI agents that can scan the web on your behalf and either bring back a canonical answer or send you personalized alerts.

The discomfort is playing out in real time. Last week, after the literary magazine *Granta* published the Commonwealth Short Story Prize–winning story “The Serpent in the Grove,” suspicious readers began to point out what they believed to be evidence of chatbot text in the story. Soon, two other Commonwealth Prize winners came under similar scrutiny, as people began running passages through AI detectors. (The Commonwealth Foundation first said in a statement that none of its prize winners had used AI, but then it issued a second statement suggesting that it is taking another look.) AI boosters celebrated the news as an example of the sophistication of current language models; skeptics viewed it as something of a slop tipping point. In a recent essay, the writer Sam Kriss described the experience of scrolling through websites right now: “The more I clicked around, the more I started to panic. There was nothing, no human voices anywhere, just thousands of versions of the same cheery demon. Am I alone out here? Something’s happened to the world; it’s all gone flimsy.”

The glut of AI writing, the detection arms race, and the debates over what constitutes appropriate use are part of the bigger questions posed by this technology. What does it mean to outsource our creativity? How long will we have the ability to discern whether something we like is human or not? Does our taste matter? If it doesn’t, then what are we even doing here?

It’s unsurprising, perhaps, that at the same time when Silicon Valley is building and breathlessly promoting these tools—self-directed agents that can accomplish complex tasks without human supervision—many of its loudest voices have grown obsessed with the idea of their own agency. In builder circles, people deemed “high-agency” sit atop the hierarchy. They are individualistic, ambitious, focused. They *just do things*. They are especially adept at marshaling the use of people and machines alike. It is implied that those with high agency are, for now, insulated from becoming replaceable or irrelevant in a time of great precarity—not yet doomed to be part of “the permanent underclass,” another Bay Area coinage for the late adopters who will be left behind. How could a person hear such language and not feel at least a little paranoid?

AI companies use the term *human-in-the-loop* to describe the relationship between

humans and AI tools in everything including chatbots and warfare. The humans perform managerial tasks: They prompt, evaluate, approve, monitor, correct. Being in the loop is meant to sound active, but the truth is, beyond the prompt, what humans are so often doing is reacting to an interaction of multiple machines. As chatbot- and AI-assisted search has outsourced web exploration to language models, this dynamic has become a primary way that humans interact with the internet. Sifting through a mix of AI- and human-generated videos, images, and text on your social-media feeds? You are passively consuming an interaction between an algorithm and things made by a computer. This is why being glued to our phone and feeds can feel so extractive and joyless, even numbing.

Recently, the online-culture researcher Aidan Walker memorably described the user experience across much of the internet as akin to the user being cuckolded by the endless scroll. People have long theorized that reliance on algorithms and the flood of bots and fake content have led to an internet that's effectively dead. Walker's theory is that the internet is not "dead" or "fake" but that models and algorithms have the bulk of agency online. Humans in the loop sit and watch, voting on short-form videos and giving feedback to the machines with every swipe. It's a bleak vision. It's also difficult to argue with.

In this system, it's not hard to see how mistrust develops, how one might suspect ulterior motives behind every piece of information they get served. *Are these really the best wireless speakers, or am I falling for SEO slop? Is this band I don't like really popular? What is popularity anymore?* Reality starts to blur. Everything goes flimsy.

The AI companies and boosters speak the language of empowerment. They're not wrong in the sense that the tools are powerful and, in many cases, quite useful. But it's tough to overstate how much these tools represent a reversal of the early promise of the weird and wild internet; of user-generated content; of stumbling upon information, people, and communities; of the crackling-static feeling of real people on the other end of the modem. In 2011, the writer Paul Ford described the internet as a "customer service medium," arguing that "humans have a fundamental need to be consulted, engaged, to exercise their knowledge (and thus power), and no other medium that came before has been able to tap into that as effectively." Ford was describing the impulse behind the creation of user-generated-content sites such as Reddit and Wikipedia but also the rise of comment sections, likes and thumbs-up reactions to posts and videos, and the very instincts that led to the rise of social media,

before it became quite so algorithmically mediated.

Revisiting that essay now, in the middle of our current agency crisis, I'm struck by the fact that Silicon Valley may have taken Ford's words to heart, but in service of the opposite outcome. It has built arguably the most impressive customer-service medium in the history of our species—an infinitely scalable, highly personalized answer machine that flatters our insecurities and mimics our idiosyncrasies. But this system automates, and even negates, the human need to be consulted. Now we consult the chatbots and they provide canonical answers. Generative AI's much-touted efficiencies also often erase the necessity of collaboration. They do the problem-solving, the heavy lifting. It is no surprise, then, that the backlash to AI, particularly the opposition to data centers, has come in the form of protests and public comments at town and city-council meetings across the country. People are taking up their agency in one of the few places they can: the physical world.

The AI companies seem to be missing all of this. To them, only the information—not the humanity—has value for their models. This is not right, but so many of us, lost in a swirl of ceaseless information, will be lulled into thinking that it is. In the new paradigm, we are not so much consulted as tasked with feeding data into the machinery. We perform our humanity, and the machine learns to mimic us. Its goal is to be better at whatever we're doing than we are.

The AI industry has ushered in an era of offshoring that's as irresistible to some people as it is disorienting to others. What's at stake is not just truth or trust, but also a sense of direction, orientation, and purpose for all of us. *Do your own research* became shorthand for the problem of the internet leading people to ivermectin hawkers and Pizzagate conspiracy theorists. But it could just as easily describe the thrill of autonomy promised by the open internet. We may miss it when we fully reckon with what comes next: a black-box machine of obscure corporate motivation that simply tells us, *Here is the answer*.

[View Discussion](#) 136

ABOUT THE AUTHOR

Charlie Warzel