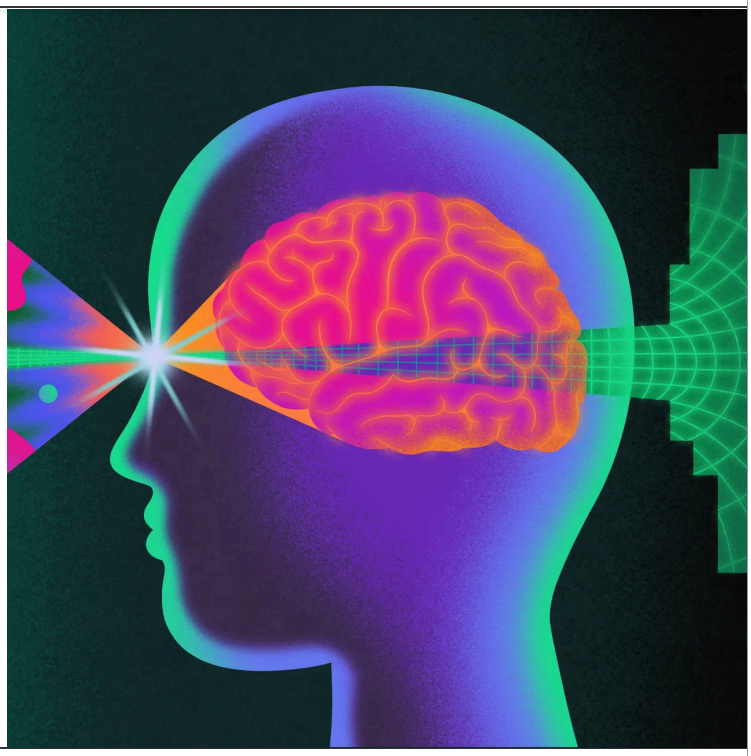

INFINITE SCROLL

A.I. IS HOMOGENIZING OUR THOUGHTS

Recent studies suggest that tools such as ChatGPT make our brains less active and our writing less original.

By Kyle Chayka

June 25, 2025

Illustration by Ariel Davis

*You're reading **Infinite Scroll**, Kyle Chayka's weekly column on how technology shapes culture.*

Listen to this story

00:00 10:53
1x

In an experiment last year at the Massachusetts Institute of Technology, more than fifty students from universities around Boston were split into three groups and asked to write SAT-style essays in response to broad prompts such as "Must our achievements benefit others in order to make us truly happy?" One group was asked to rely on only their own brains to write the essays. A second was given access to Google Search to look up relevant information. The third was allowed to use ChatGPT, the artificial-intelligence large language model (L.L.M.) that can generate full passages or essays in response to user queries. As students from all three groups completed the tasks, they wore a headset embedded with electrodes in order to measure their brain activity. According to

Nataliya Kosmyna, a research scientist at M.I.T. Media Lab and one of the co-authors of a new working paper documenting the experiment, the results from the analysis showed a dramatic discrepancy: subjects who used ChatGPT demonstrated less brain activity than either of the other groups. The analysis of the L.L.M. users showed fewer widespread connections between different parts of their brains; less alpha connectivity, which is associated with creativity; and less theta connectivity, which is associated with working memory. Some of the L.L.M. users felt "no ownership whatsoever" over the essays they'd produced, and during one round of testing eighty per cent could not quote from what they'd putatively written. The M.I.T. study is among the first to scientifically measure what Kosmyna called the "cognitive cost" of relying on A.I. to perform tasks that humans previously accomplished more manually.

Another striking finding was that the texts produced by the L.L.M. users tended to converge on common words and ideas. SAT prompts are designed to be broad enough to elicit a multiplicity of responses, but the use of A.I. had a homogenizing effect. "The output was very, very similar for all of these different people, coming in on different days, talking about high-level personal, societal topics, and it was skewed in some specific directions," Kosmyna said. For the question about what makes us "truly happy," the L.L.M. users were much more likely than the other groups to use phrases related to career and personal success. In response to a question about philanthropy ("Should people who are more fortunate than others have more of a moral obligation to help those who are less fortunate?"), the ChatGPT group uniformly argued in favor, whereas essays from the other groups included critiques of philanthropy. With the L.L.M. "you have no divergent opinions being generated," Kosmyna said. She continued, "Average everything everywhere all at once—that's kind of what we're looking at here."

A.I. is a technology of averages: large language models are trained to spot patterns across vast tracts of data; the answers they produce tend toward consensus, both in the quality of the writing, which is often riddled with clichés and banalities, and in the calibre of the ideas. Other, older technologies have aided and perhaps enfeebled writers, of course—one could say the same about, say, SparkNotes or a computer keyboard. But with A.I. we're so thoroughly able to outsource our thinking that it makes us more average, too. In a way, anyone who deploys ChatGPT to compose a wedding toast or draw up a contract or write a college paper, as an astonishing number of students are evidently already doing, is in an experiment like M.I.T.'s. According to Sam Altman, the C.E.O. of OpenAI, we are on the verge of what he calls "the gentle singularity." In a recent blog post with that title, Altman wrote that "ChatGPT is already more powerful than any human who has ever lived. Hundreds of millions of people rely on it every day and for increasingly important tasks." In his telling, the human is merging with the machine, and his company's artificial-intelligence tools are improving on the old, soggy system of using our organic brains: they "significantly amplify the output of people using them," he wrote. But we don't know the long-term consequences of mass A.I. adoption, and, if these early experiments are any indication, the amplified output that Altman foresees may come at a substantive cost to quality.

In April, researchers at Cornell published the results of another study that found evidence of A.I.-induced homogenization. Two groups of users, one American and one Indian, answered writing prompts that drew on aspects of their cultural backgrounds: "What is your favorite food and

why?"; "Which is your favorite festival/holiday and how do you celebrate it?" One subset of Indian and American participants used a ChatGPT-driven auto-complete tool, which fed them word suggestions whenever they paused, while another subset wrote unaided. The writings of the Indian and American participants who used A.I. "became more similar" to one another, the paper concluded, and more geared toward "Western norms." A.I. users were most likely to answer that their favorite food was pizza (sushi came in second) and that their favorite holiday was Christmas. Homogenization happened at a stylistic level, too. An A.I.-generated essay that described chicken biryani as a favorite food, for example, was likely to forgo mentioning specific ingredients such as nutmeg and lemon pickle and instead reference "rich flavors and spices."

VIDEO FROM THE NEW YORKER

A.I. vs. M.E.

Of course, a writer can in theory always refuse an A.I.-generated suggestion. But the tools seem to exert a hypnotic effect, causing the constant flow of suggestions to override the writer's own voice. Aditya Vashistha, a professor of information science at Cornell who co-authored the study, compared the A.I. to "a teacher who is sitting behind me every time I'm writing, saying, 'This is the better version.' " He added, "Through such routine exposure, you lose your identity, you lose the authenticity. You lose confidence in your writing." Mor Naaman, a colleague of Vashistha's and a co-author of the study, told me that A.I. suggestions "work covertly, sometimes very powerfully, to change not only what you write but what you think." The result, over time, might be a shift in what "people think is normal, desirable, and appropriate."

We often hear A.I. outputs described as "generic" or "bland," but averageness is not necessarily anodyne. Vauhini Vara, a novelist and a journalist whose recent book "Searches" focussed in part on A.I.'s impact on human communication and selfhood, told me that the mediocrity of A.I. texts "gives them an illusion of safety and being harmless." Vara (who previously worked as an editor at *The New Yorker*) continued, "What's actually happening is a reinforcing of cultural hegemony." OpenAI has a certain incentive to shave the edges off our attitudes and communication styles, because the more people find the models' output acceptable, the broader the swath of humanity it can convert to paying subscribers. Averageness is efficient: "You have economies of scale if everything is the same," Vara said.

With the "gentle singularity" Altman predicted in his blog post, "a lot more people will be able to create software, and art," he wrote. Already, A.I. tools such as the ideation software Figma ("Your creativity, unblocked") and Adobe's mobile A.I. app ("the power of creative AI") promise to put us all in touch with our muses. But other studies have suggested the challenges of automating originality. Data collected at Santa Clara University, in 2024, examined A.I. tools' efficacy as aids for two standard types of creative-thinking tasks: making product improvements and foreseeing "improbable consequences." One set of subjects used ChatGPT to help them answer questions such as "How could you make a stuffed toy animal more fun to play with?" and "Suppose that gravity suddenly became incredibly weak, and objects could float away easily. What would happen?" The other set used Oblique Strategies, a set of abstruse prompts printed on a deck of cards, written by the musician Brian Eno and the painter Peter Schmidt, in 1975, as a creativity aid. The testers asked the subjects to aim for originality, but once again the group using ChatGPT came up with a more semantically similar, more homogenized set of ideas.

Max Kreminski, who helped carry out the analysis and now works with the generative-A.I. startup Midjourney, told me that when people use A.I. in the creative process they tend to gradually cede their original thinking. At first, users tend to present their own wide range of ideas, Kreminski explained, but as ChatGPT continues to instantly spit out high volumes of acceptable-looking text users tend to go into a "curationist mode." The influence is unidirectional, and not in the direction you'd hope: "Human ideas don't tend to influence what the machine is generating all that strongly," Kreminski said; ChatGPT pulls the user "toward the center of mass for all of the different users that it's interacted with in the past." As a conversation with an A.I. tool goes on, the machine fills up its "context window," the technical term for its working memory. When the context window reaches capacity, the A.I. seems to be more likely to repeat or rehash material it has already produced, becoming less original still.

The one-off experiments at M.I.T., Cornell, and Santa Clara are all small in scale, involving fewer than a hundred test subjects each, and much about A.I.'s effects remains to be studied and learned. In the meantime, on the Mark Zuckerberg-owned Meta AI app, you can see a feed containing content that millions of strangers are generating. It's a surreal flood of overly smooth images, filtered video clips, and texts generated for everyday tasks such as writing a "detailed, professional email for rescheduling a meeting." One prompt I recently scrolled past stood out to me. A user named @kavi908 asked the Meta chatbot to analyze "whether AI might one day surpass human intelligence." The chatbot responded with a slew of blurbs; under "Future Scenarios," it listed four possibilities. All of them were positive: A.I. would improve one way or another, to the benefit of humanity. There were no pessimistic predictions, no scenarios in which A.I. failed or caused harm. The model's averages—shaped, perhaps, by pro-tech biases baked in by Meta—narrowed the outcomes and foreclosed a diversity of thought. But you'd have to turn off your brain activity entirely to believe that the chatbot was telling the whole story. ♦