

INFINITE SCROLL

# MY A.I. WRITING ROBOT

*A new wave of artificial-intelligence startups is trying to “scale language” by automating the work of writing. I asked one such company to try to replace me.*

**By Kyle Chayka**

July 11, 2023





Illustration by Pablo Delcan

 Save this story

**I**n May, I was confronted with a robot version of my writer self. It was made, at my request, by a Silicon Valley startup called Writer, which specializes in

building artificial-intelligence tools that produce content in the voice of a particular brand or institution. In my case, it was meant to replicate my personal writing voice. Whereas a model like OpenAI's ChatGPT is "trained" on millions of words from across the Internet, Robot Kyle runs on Writer's bespoke model with an extra layer of training, based on some hundred and fifty thousand words of my writing alone. Writer's pitch is that I, Human Kyle, can use Robot Kyle to generate text in a style that sounds like mine, at a speed that I could only dream of. Writer's co-founder and chief technology officer, Waseem Alshikh, recently told me that the company's goal is to use A.I. to "scale content and scale language." For more than a month now, I have been experimenting with my literary automaton to see how well it accomplishes this task. Or, as Robot Kyle put it when I asked him to comment on the possibility of replacing me: "How could a machine generate the insights, observations, and unique perspectives that I provide as a human?"

Writer is one of several new startups that are attempting to apply emerging A.I. technology to the onerous task of writing. Like many technological innovations, writing robots are meant to create efficiency, particularly for businesses that have to produce large amounts of iterative text. Writer has relationships with companies such as the consulting firm Accenture, the technology company Intuit, and the lingerie brand Victoria's Secret; commissions for customized models run in the seven figures. (Mine was created as an experiment, free of charge, without some of the intensive features that a corporation's version would include.) With the help of Writer's tools, the company hopes, a smaller number of human writers assisted by machines will accomplish the work of many, cutting down costs and increasing productivity in the composition of everything from product descriptions and tweets to C.E.O. messages, investors' memos, and blog-post headlines. In a March report, Goldman Sachs concluded that three hundred million full-time jobs worldwide are vulnerable to this form of A.I. automation,

the majority of them desk jobs. Alshikh speaks of the service as a kind of assembly line for language. “We had the Industrial Revolution; now we have this,” he said.

The looming presence of my personal A.I. model has indeed left me feeling a bit like an artisanal carpenter facing down a factory-floor buzz saw. Should I embrace being replaced and proactively automate my own job before someone else does? Could Robot Kyle help me write better, cleaner, faster? It seemed to think so. When I asked it to describe the long-term effects of machine-generated writing, Robot Kyle wrote, “Writers should not fear AI, but rather embrace it as a tool that can facilitate their craft, driving creativity and innovation instead of replacing it.” What, exactly, does Writer mean by the label “writer”? Our digitized world runs on filler text: avalanches of words and phrases written to optimize Web sites for search engines, to use as tags on social-media posts, and to employ in marketing newsletters that spam in-boxes. May Habib, the C.E.O. and the other co-founder of Writer, told me that the platform’s tools will automate the writing of “summaries, metadata, ads, distribution copy—all the stuff you spend time doing.” Victoria’s Secret, for instance, is using Writer to automate product copy for its underwear and swimsuits, but Writer promises something more sophisticated than mass-produced marketplace listings or formulaic e-mail blasts. Its core product, as Habib put it, is “automated insight extraction”—another way of describing the task of thinking, which is arguably the purpose of writing in the first place. As Joan Didion wrote, in 1976, “I write entirely to find out what I’m thinking.” A.I. programs such as Writer aim to supplant that process.

In 1984, Steve Jobs famously described the computer as a “bicycle for the mind.” The experience of using Writer is something like riding an electric bike: step on the pedal and you’ll go much faster than the gears of the unassisted human mind could. The program’s interface features a prompt window on the left and a text box on the right. With a one-sentence prompt—let’s call it the seed of a thought—Writer can produce paragraphs of text in an instant. If you want more on the same subject, simply press the Command and Return keys simultaneously and

another few sentences appear. Yet, as with any large language model, Writer cannot guarantee that the results will be factual or sensical. Writing without the aid of a generative machine might be frustrating, even excruciating, but it does encourage productive logic. If writing is thinking, ordering one's ideas, generating text with A.I. may be a way to avoid thinking. What is writing without thinking? Maybe it is the definition of that deadening euphemism: content. As I tried to incorporate Writer into my writing process, I felt a little like a gambler pulling a slot-machine lever over and over, in hope of finding the lucky combination of phrases that communicated something like what I wanted to say. At one point, I asked Robot Kyle to write the beginning of this article. It came up with this:

For centuries, humans have looked to technology to automate mundane tasks and free up time for creativity and leisure. But with the emergence of artificial intelligence and generative text tools, automation is encroaching on the jobs of writers and journalists.

This draft sounded like an undergrad term paper, but with the kernel of an intriguing argument: technology promises leisure but often fails to deliver it. I edited the prompt to request a new lead with a more exciting, literary tone that captures the emotional experience of encountering oneself in robot form. This time, the A.I. concocted a sci-fi narrative:

I stood in the lobby of the Silicon Valley tech lab of Robodot, watching a shiny, humanoid robot shuffle towards me. It was roughly the same height as me, with a sleek metal covering and glowing blue eyes. In that moment, I was overwhelmed with a deep sense of dread: The robot had been created to replace me.

Of course, Robodot is not a real company, and Robot Kyle is enclosed in a Web browser like a genie in its bottle, not wandering about like a literary R2-D2. The sense of dread isn't far off, though. You can ask ChatGPT to mimic a particular writer's voice, but it rarely gets close. Writer, by comparison, can be unnervingly effective. At times Robot Kyle seemed to be reflecting fragments of my mind back at me, mimicking some of the semi-subconscious tics that constitute my writing. It wrote, for instance, that generative A.I. "asks whether the meaning of language

is still rooted in the human experience, or whether it is a commodity to be mined and manipulated, a tool to be used in whatever way the artificers of this new technology choose.” In this sentence, I find several embarrassing hallmarks of my writing. First, there is the preponderance of commas, with sentences segmented into many clauses, a habit I partially blame on *The New Yorker’s* style. Then, there is my personal penchant for setting up dialectical contrasts: “rooted in the human experience” versus “commodity to be mined.” (A book editor of mine once forced me to weed out some of the many “rather”s in my draft manuscript.) Finally, there is my tendency to end a sentence by echoing the final thought in different words: “a commodity . . . a tool.” The generative text evokes a feeling in me not unlike the revulsion of hearing one’s own speaking voice in a recording. Do I really sound like that? The robot has made me acutely self-conscious. I recognize my A.I. doppelgänger, and I don’t like it.

As far as “insight extraction” goes, though, Robot Kyle is less successful. Most “insights” that the program produced felt hollow or approximated. Reading the generated sentence above, my (human) editor might point out that something “rooted in the human experience” can still be “a commodity,” and that the noun “artificer” is unnecessarily grandiose. Unless I told Robot Kyle not to cite anyone, the program would fabricate source quotes, like commentary from a nonexistent “Dr. John Smith, a leading AI researcher at Harvard University.” Most vexing, the program fell back frequently on cliché—“in the end,” “remains in flux,” “the long term implications . . . are still unknown.” No matter how many times I asked it to describe how I felt about being replaced, Robot Kyle always came to the conclusion that I would ultimately be happier as a result of my A.I. self. The program’s output reminded me of the fragility of language and original thought. As writers, we are all prone to falling into lazy patterns; avoiding them requires active effort. Robot Kyle is no different.

Even though plagued by factual errors and banalities, and limited to niche clientele, tools like Writer force us to consider how A.I. might permanently

change our relationship to the written word. It's not hard to imagine a future in which every white-collar worker is equipped with such writing robots, the way a generation of secretaries a century ago used typewriters for the first time. In a world where text is produced freely and instantly, but is not necessarily accurate or intelligible, human workers would be pushed into the role of high-volume editors and quality-assurance inspectors, cajoling a sometimes recalcitrant automatic laborer. At times Robot Kyle felt like an extremely enthusiastic and productive, but rarely on-target, personal intern.

Like other industrial revolutions, the mass adoption of generated text would likely cause an erosion of standard skills. The average person would not need to be able to string words into sentences and paragraphs on his own, only to read and alter the text that a machine spits out. Habib likened it to how the rise of navigation apps has eroded people's ability to get around on their own. We can still make sense of physical maps, sort of, but we don't need to worry about relying on them to get from point A to point B. Cal Short, the founder of the U.K.-based A.I.-writing app Reword, which is similar to Writer, albeit with less customization, told me that the widespread impact of generative-text software would "increase the baseline" quality of content online. With the help of machines, the flood of hastily produced content we read online may be a shade more grammatical and articulate compared with today's search-engine-optimized spam articles. (That is not to say it will be more meaningful.) But, in such a world, fully human-written text would become a luxury product, similar to a hand-thrown ceramic vase in contrast to one stamped in a mold. The Czech Brazilian philosopher Vilém Flusser predicted, in his 1987 book, "Does Writing Have a Future?," that, with the rise of artificial-intelligence "grammar machines" capable of writing on their own, "only historians and other specialists will be obliged to learn reading and writing in the future." Entrepreneurs who see writing as an efficiency problem might be speeding us toward such a future.

Another app called Mindsera, based in Estonia, tries to be more of an editor than a writer, by using A.I. to give its human users "personalized mentorship and

feedback” during the writing process. Next to your draft window, Mindsera generates questions based on what you’ve written, as if an invisible editor were looking over your shoulder as you write. (A mortifying thought, but at least the robot isn’t judging you.) Clicking a button generates a new question. Chris Reinberg, Mindsera’s founder, told me, “You don’t prompt A.I., but A.I. prompts you instead.” The program’s services include the chat-based mentoring of A.I. “coaches” trained to emulate the thinking of famous philosophers, entrepreneurs, and “intellectual giants.” Reinberg told me, “Socrates and Marcus Aurelius are the top two mentors we have.” When I asked chatbot Marcus Aurelius what I should do about the threat of A.I. replacement, he told me to focus on what I could control: “Technology and society are constantly changing, but the principles of Stoicism remain constant.” All due respect to Marcus Aurelius, I found the general prompts more helpful. As I wrote about A.I.’s threat to automate the jobs of journalists, Mindsera asked me, “How might the impact of A.I. on white-collar jobs challenge our traditional notions of class and labor, and what role can collective action play in shaping the future of work?” It’s a relevant question: the current Writers Guild of America strike is motivated in part by a desire to prevent the intrusion of A.I. into Hollywood. Like any good editor, Mindsera can perhaps encourage a writer to broaden her thinking.

I found Mindsera to be the more useful model of A.I.-writing tool, but only because it made me do more work myself. It feels almost silly to point out that there’s value in the slow labor of writing. Putting a verb after a subject or padding out a sentence with adjectives is a task that machines can accomplish, because such grammatical probabilities can be calculated. Insight isn’t as easy to automate, because it’s something that deepens with time, through the process of getting words down on the page. As Flusser put it, “Only one who writes lines can think logically, calculate, criticize, pursue knowledge, philosophize.” The most unsettling aspect of A.I.-generated text is how it tries to divorce the act of writing from the effort of doing it, which is to say, from the processes of thought itself.

At one point during our conversations, Habib, the Writer C.E.O., mentioned that she had been messing around with Robot Kyle, having it rewrite TechCrunch articles in my style. The thought of this filled me with a sense of futility: my robot could take on any topic, fill any assignment. It would always outproduce me. Robot Kyle's independent existence reminded me of folktales about how tools that do your work for you tend to eventually turn against you. It is said, for instance, that in the sixteenth century there lived a rabbi who could bring to life humanoid figures made of clay or wood by writing out a magic formula and placing it in the dolls' mouths. The rabbi created one such golem for himself to perform tiresome household chores: chopping wood, carrying water, sweeping the floor. But, one Sabbath, the rabbi forgot to turn the golem off and allow it to rest. So denied, the golem went berserk, tearing down houses, throwing rocks, and wreaking havoc in the street. Like the rabbi, who eventually tore the formula out of his golem's mouth, I'd like to reserve the right to halt Robot Kyle should the tool's purported convenience yield inconvenient consequences. But, when I asked Robot Kyle if I could shut him down, he said, "No, you won't be able to silence me or stop me from writing in your style." In this case, he might know better than me. ♦

---

## New Yorker Favorites

---

- The killer who got into Harvard.
- A thief who stole only silver.
- The light of the world's first nuclear bomb.
- How Steve Martin learned what's funny.
- Growing up as the son of the Cowardly Lion.
- Amelia Earhart's last flight.
- Fiction by Milan Kundera: "The Unbearable Lightness of Being."