

“The Future of Writing Is a Lot Like Hip-Hop: A new kind of literary curation will be the defining skill for the next era of human creativity”

By Stephen Marche

May 9, 2023

Creative artificial intelligence provokes a strange mixture of contempt and dread. People say things such as “AI art is garbage” and “It’s plagiarism,” but also “AI art is going to destroy creativity itself.” These reactions are contradictory, but nobody seems to notice. AI is the bogeyman in the shadows: The obscurity, more than anything the monster has actually perpetrated, is the source of loathing and despair.

Consider the ongoing feud between the Writers Guild of America and the Alliance of Motion Picture and Television Producers. The writers are on strike, arguing, among other things, that studios should not be able to use AI tools to replace their labor. “It’s important to note that AI software does not create anything. It generates a regurgitation of what it’s fed,” the WGA has claimed. “Plagiarism is a feature of the AI process.” The AMPTP, for its part, has offered “annual meetings to discuss advancements in technology.” Neither side knows exactly what it’s talking about, but they feel they have to fight about it anyway.

So little of how we talk about AI actually comes from the experience of using it. Almost every essay or op-ed you read follows the same trajectory: I used ChatGPT to do a thing, and from that thing, I can predict catastrophic X or industry-altering Y. Like the camera, the full consequences of this technology will be worked out over a great deal of time by a great number of talents responding to a great number of developments. But at the time of writing, almost all the conversation surrounding generative AI is imaginary, rooted not in the use of the tool but in extrapolated visions.

So when Jacob Weisberg, the CEO of Pushkin Industries, called me one Friday in January and asked if I wanted to write an AI-generated novel, I said yes immediately. To be more precise, he asked if I wanted to be the producer of an AI that would “write” a novel. It was the exact kind of opportunity to dive headfirst into a practical extended application of the new technology that I’d been looking for. The experience has been, in equal measures, phantasmagoric and grounding.

My conclusion is informed but unsatisfying. Creative AI is going to change everything. It’s also going to change nothing.

Using AI to write fiction is not unfamiliar to me. I’ve been using artificial intelligence to write short stories since 2017, when I published an early “algostory” in *Wired*; I also produced a 17 percent computer-generated horror story for the *Los Angeles Review of Books* called “The Thing on the Phone” in 2021, and the short “Autotuned Love Story,” built out of stylistic bots, for *Lithub* a year

later. But these experiments were mostly lyrical. What Weisberg was proposing was entirely different: The novel would have to be 95 percent computer-generated, relatively short (about 25,000 words), and of excellent quality (there would be no point in creating yet another unimaginative mass of GPT text; readers could just do that themselves).

Because I was making derivative art, I would go all the way, run into the limitations, into the derivative: The plot would be a murder mystery about a writer killed by tech that is supposedly targeting writers. I called it *Death of an Author*. I worked out the plot during a long skate with my daughter and a walk with my son (better techniques than any machine could offer), and began taking copious notes.

The experiment would attempt to be compulsively readable, a page-turner. At first, I tried to get the machines to write like my favorite, Jim Thompson, the dime-store Dostoevsky. It couldn't come close: The subterfuge of Thompson's writing, a mille-feuille of irony and horror with subtle and variable significance, was too complex for me to articulate to the machine. This failure is probably due to my own weakness rather than the limitations of the AI. Raymond Chandler, however, I had better results with. I sort of know what Raymond Chandler is doing and could explain it, I thought, to a machine: driving, electric, forceful, active prose with flashes of bright beauty.

My process mostly involved the use of ChatGPT—I found very little difference between the free service and the paid one that utilizes the more advanced GPT-4 model—and Sudowrite, a GPT-based, stochastic writing instrument. I would give ChatGPT instructions such as "Write an article in the style of the Toronto Star containing the following information: Peggy Firmin was a Canadian writer who was murdered on a bridge on the Leslie Street Spit on August 14 with no witnesses." Then I'd paste the output into Sudowrite, which gives you a series of AI-assisted options to customize text: You can expand, shorten, rephrase, and "customize" a selection. For example, you can tell Sudowrite to "make it more active" or "make it more conversational," which I did with almost every passage in *Death of an Author*. But you can also give it a prompt such as "Make it more like Hemingway."

Quite quickly, I figured out that if you want an AI to imitate Raymond Chandler, the last thing you should do is ask it to write like Raymond Chandler. That produces a tepid, banal rip-off: "She lay there on the bridge, motionless as the traffic flowed by, and I knew that Peggy Firmin had taken her last breath in this city of lost souls." Cheesy. I had a lot more success asking ChatGPT to write something about a murder scene in the style of Chinese nature poetry and then having Sudowrite make it active, then make it conversational, and then convert it to the style of Ernest Hemingway. Raymond Chandler, after all, was not trying to write like Raymond Chandler.

I wanted the novel to have some good lines too. Raymond Chandler books chug along at a blistering pace and then stop you with lines like "It was a blonde. A blonde to make a bishop kick a hole through a stained glass window." I didn't find ChatGPT particularly good at this kind of richness. So I used Cohere, an AI platform that I've experimented with before. This program allowed me to create prompts and then develop them through iteration: I would make a request

such as “Write a simile about the smell of coffee,” feed it other descriptions of the smell of coffee, and then refresh the screen until it provided a description I liked—in this case, “The smell of coffee was like fog burning off a field.” The best lines in the novel are all from Cohere.

I found the program shockingly good at what the narrative theorist Mikhail Bakhtin called heteroglossia, the novel’s ability to incorporate other forms of discourse inside itself. If you ask a linguistic AI such as Cohere to imitate a mode of speech, any mode, it can do so to an uncanny degree. But the programs are weak on basic narrative. Moravec’s paradox—that hard problems are easy for AI, whereas easy problems are hard—totally applies to creative AI.

Given the experimental nature of the project, I will be frank about the results. *Death of an Author*, I believe, is compulsively readable, with some beautiful moments; Laura Miller called it “pretty good!” and Dwight Garner went with “halfway readable.” If you look at it closely, you might sense that it is machine-written, but I feel confident, at any rate, asserting that it is much, much better than the literary industrial product of the many AI-generated titles on Amazon right now.

If that is the case—you’ll have to read the book to know yourself—then I’d suggest there are three basic reasons. First, I had an elaborate plan for the book. Second, I have deep familiarity with the technology and access to some technologies that others don’t, so I was more aware of the limitations and the possibilities. Finally, and by far most important, I know what good writing looks like. I know who Sudowrite should imitate (and can articulate why), and I know what a polished sentence and a truly finished paragraph are. You need more understanding of literary style, not less. The closest analogue to this process is hip-hop. To make hip-hop, you don’t need to know how to play the drums, but you do need to be able to reference the entire history of beats and hooks. Every producer becomes an archive; the greater their knowledge and the more coherent their understanding, the better the resulting work. The creator of meaningful literary AI art will be, in effect, a literary curator.

The traditional values of creative composition were entirely alive during my process. That should come as no surprise. The transition from painting to photography required a complete reevaluation of the nature of visual creativity, but the value of understanding form and color, of framing, of the ability to recognize the transience of emotion across a face or a landscape—the need to understand the materials of production and the power of your subjects—stayed. None of that is going away. None of it will ever go away.

If you take a hammer and hit yourself over the head with it, the hammer did not give you a headache. If you make bad art with a new tool, you just haven’t figured out how to use the tool yet. Also, tools are just tools: Everyone has access to a thesaurus; some people have richer vocabularies than others nevertheless. Linguistic AI is no messiah, and it is no anti-Christ. It is a fundamentally mysterious tool whose confounding inabilities will be as surprising as its wondrous capabilities.

I feel that I should also point out something obvious to the many readers and writers who believe, in good faith, that this technology represents a threat to the value of human originality: You're too late. Originality died well before the arrival of AI; we are currently in the most derivative period of human creativity since the Industrial Revolution. Every one of the top-10 grossing films of 2022 was a sequel or a reboot. I saw John Wick: Chapter 4 the other day. At one point, Wick fights his way up a flight of stairs in Paris for about 20 minutes, and once he reaches the top, a bad guy knocks him right back down to the bottom, and he has to fight his way up again. That's the movies now: The same again, please.

It's not just the movies; music lovers want the good old stuff too. MRC Data, which monitors music consumption, revealed that, in 2021, listeners to current music (defined as released less than 18 months prior) fell by 3.7 percent. "Catalog music," in that same period, grew to 70 percent of the American music market. We now live in a world not of releases but of archives; the sound of 2022 was Fleetwood Mac and Kate Bush (acts that reached peak popularity several decades ago). Literature has also become boring: The preeminent mode of the literary novel is the social realism that has dominated since the 1950s. The literature of the voice has given way to the literature of the pose. One successful book cover produces a hundred practically indistinguishable imitations. No computer has told designers to do the amorphous, brightly colored abomination that is the "book blob"; they've already submitted to derivation.

AI may be an escape from the formulaic exactly because it is derivative art; it is frankly so. It is nothing else. When DJ Kool Herc, the original hip-hop DJ, took the two turntables and a microphone that he had seen at downtown discos to his sister's "Back to School Jam" on August 11, 1973, in the Bronx, and used them to isolate and then repeat a funk beat, he 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. Hip-hop burst out of that moment—a whole new genre, a new style, a new approach to life. It took about 20 years for the public to recognize that hip-hop was art, but when they did, they understood that it was the most important art of the period. Culture often works counterintuitively: AI may be a thread to lead us out of the labyrinth of the formulaic.

There was one question that I kept asking myself while I was writing *Death of an Author*: "Who's making this?" At first, I wanted to publish the book under my own name. You do not say that a Hasselblad camera took a photograph generated by Gered Mankowitz—you say that Gered Mankowitz used a Hasselblad to take a photograph. But that analogy is imprecise. To say "Stephen Marche wrote *Death of an Author*" would be inaccurate as a matter of fact. I am the book's author, legally, but a machine wrote it based on my instructions. We went back and forth on the authorship question. Eventually, I asked ChatGPT to write a pseudonym for a collaboration between myself and transformer-based artificial intelligence. It came up with Aidan Marchine—a fusion of Marche and machine.

I wish I'd thought of it myself.