A CRITIC AT LARGE

ROBOT LIT

The prehistory of A.I. slop.

BY JILL LEPORE

In 1962, a programmer at Librascope, a California-based defense contractor, announced that "a computer can be programmed to write meaningful and relevant sentences in proper English." At Librascope's Laboratory for Automata Research, in Glendale, he'd started out by feeding into his computer—the vacuum-tube LGP-30—a vocabulary of thirty-five hundred words and a repertoire of a hundred and twenty-eight sentence patterns, and told it to do, more or less, what humans did in the nineteen-nineties when they stuck Magnetic Poetry on the doors of their refrigerators. And behold! "Broccoli is often blind," the LGP-30 tapped out on its typewriter, and "Communism is more porcelain than albino gold." The engineer decided to set this machine-generated text as free verse:

> Was Milo mewling thrilling radishes?
> So, our anchovies are sad but green.

He called his program the Auto-Beatnik, cunningly deploying Kerouac, Ginsberg, and Burroughs as cover for this bilge. The ploy occasionally paid off; the London *Daily Mirror* described the Auto-Beatnik's poetry as "better than most of the stuff that gets published in avant garde magazines."

The literary début of the Auto-Beatnik, a machine that could compose five thousand poems in an hour or so, caught the attention of *Time*, *Life*, and the *Times*. In an inversion of the more common critical reception of an emerging artist, this new writer's poetry was often noticed but seldom admired, notwithstanding the *Daily Mirror's* snide enthusiasm. In 1963, a story in the *Observer* about automated poetry ran with a cartoon of a guy feeding a slip of paper reading "ART" into a computer that, at the other end, spat out a paper reading "TRA." The "SH," I guess, was implied: in went art, out came trash.

Lately, this kind of junk has become known as A.I. slop—"slop" was Merriam-Webster's 2025 Word of the Year—and it's everywhere, gumming up the works, slowing down traffic, and making a god-awful mess. It brings to mind the time, in 1919, that a tank in Boston containing nearly two and a half million gallons of molasses burst, and a fearsome wave of syrup reportedly fifty feet high and travelling at thirty-five miles an hour (faster than you'd expect, really) flooded the city. The cleanup of the Great Molasses Flood took weeks, and, for months, everywhere that anyone had tracked molasses, including underground subway platforms, was still tacky. Even years later, on hot days, the North End smelled like a gingerbread house.

Machine-generated writing, though it doesn't smell as sweet, has something of molasses's smothering stickiness. One way to think about the internet is that it's an attempt to archive nearly everything ever written by anyone who ever lived. Recently, more and more new writing online is being produced by bots, during this, the Great A.I.-Slop Flood. Ante-ChatGPT, more than ninety-eight per cent of all English-language articles being published on the internet were written by humans. By the fall of 2024, machines were writing around half of such articles, according to the digital-marketing agency Graphite, which, far from taking umbrage at the usurpers, recommends using A.I. to help run your ad campaigns. And why not? In one blind test, people found A.I.-generated advertisements to be "of higher quality" than ads made by humans.

And that's not counting social media or e-mail or all the robot-written rubbish that comes your way by text or voice mail or pop-up customer-service chats. YouTube is overrun with slop. Reddit is caked in it. Much of Facebook is nothing *but* slop. The literary critic Matthew Kirschenbaum warns of a coming "textpocalypse" that will render the words you're reading right now—this word, and this one—relics your grandchildren will frame on a wall, a daguerreotype, a needlepoint sampler. "Like the prized pen strokes of a calligrapher, a human document online could become a rarity to be curated, protected, and preserved," Kirschenbaum writes. Can the textpocalypse be stopped? "Rest assured 2026 will be the beginning of AI slop purge," *Forbes* promised, sloppily, at the start of this year. This was hardly reassuring. My anchovies are still sad.

The idea of mechanically produced prose or poetry is not especially new. Eighteenth-century letter-writing manuals provided fill-in-the-blank templates, because many types of correspondence are set forms: letters of condolence, say, or letters of recommendation. Anxiety about machines replacing humans as writers, and replacing good writing with bad, is also older than you might think. Mid-nineteenth-century commentators, overwhelmed by the era's flood of cheap printed material, especially periodicals and novels, imagined a "New Magazine Machine"

SOURCE PHOTOGRAPH FROM GETTY; OPPOSITE: PIERRE BUTTIN

Long before ChatGPT, computers were already producing love letters, fake poems, cheap plots, and high-grade nonsense.

PHOTO ILLUSTRATION BY JACK SMYTH

SOURCE PHOTOGRAPH FROM GETTY; OPPOSITE: PIERRE BUTTIN

that could spit out cheap pulps, and a "Book-Making Machine," a literary successor to Charles Babbage's Analytical Engine.

Random-story generators are even more ancient: that, after all, is what tarot cards are. (You can read about some of these antecedents in Dennis Yi Tenen's breezy 2024 book, "Literary Theory for Robots: How Computers Learned to Write," an introduction that, despite its title, isn't really a work of literary theory but instead engages in "patiently assembling the modern chatbot from parts found on the workbench of history.") As industrialization advanced, the factory replaced the wheel of fortune as a metaphor for how things happen in the world. A 1912 writing guide, "The Fiction Factory," advised, "A writer is neither better nor worse than any other man who happens to be in trade. He is a manufacturer. After gathering his raw product, he puts it through the mill of his imagination." This only accelerated with the factory that was Hollywood. In a 1919 writing guide, "Ten Million Photoplay Plots," a grifter named Wycliffe A. Hill told would-be screenwriters that there are thirty-seven possible story lines that can be combined with a measurable number of characters, situations, and subplots to produce the mathematically precise total of 10,494,360 plots. After the coining of the word "robot," in 1920, in the internationally popular Czech play "R.U.R.," and the attendant cultural fascination with mechanical men, Hill published a follow-up manual in 1931 that included what he called the Plot Robot. As an ad for it in *Modern Mechanics* promised:

> Formerly robots were merely mechanical devices that could perform a variety of stunts under the guidance of a human being, but now a robot has made its appearance that thinks, has a soul of a kind, creative imagination, and other qualities necessary for writing a modern stereotyped short story. . . . Now if you want to become a successful author simply obtain a robot and put it to work.

In fact, there was no robot. If you bought the book, you found out that the Plot Robot was a cardboard number wheel. This grift is still going. These days, you can buy Writing Dice to help you with your novel: Nine dice! "Thousands of combinations, you'll never fear the blank page again!"

Actual robot writing dates to 1953, when the mathematician Christopher Strachey, a nephew of the writer Lytton Strachey, he of the Bloomsbury group, wrote a computer program that could generate love letters like this one:

> Honey Dear
> My sympathetic affection beautifully attracts your affectionate enthusiasm. You are my loving adoration: my breathless adoration. . . .
> Yours wistfully
> M. U. C.

Think of it as Mad Libs before there was Mad Libs. Strachey instructed M.U.C., the Manchester University Computer, to fill in the blanks in template sentences by drawing randomly from a list of words identified by their parts of speech: "*My* — (adj.) — (noun) — (adv.) — (verb) *your* — (adj.) — (noun)." He then posted the letters on campus. Among Strachey's motivations for building a cyber Cyrano was to poke fun at credulous reporters who described computers as "thinking machines." His program, he insisted, was "almost childishly simple." Because Strachey was thought to have been gay, scholars have read the letters as making fun of straight romance. Or, I'd have said, any romance. Incontestably, love letters are, very often, slobbering slop.

Lest Strachey's epistles seem antiquated compared with the stuff that comes your way these days, I might mention that, while I was writing this essay, a writer friend texted me an A.I. e-mail she'd just received that purported to be from a British novelist: "Happy weekend, Elise! Quick bulldozer boost (spam-free!) your toy-truck kid spark + 28-book marathon inspire; let's swap suspense secrets 15 mins this week? When will work reply today? Warmly, Alice Feeney." DEAR HONEY-DEW YOU ARE MY GREATEST WHISKERS MY UTTER MOONBEAM.

Strachey's work is the starting point for an engrossing collection, "Output: An Anthology of Computer-Generated Text, 1953-2023," edited by Lillian-Yvonne Bertram and Nick Montfort, but they do not note that 1953 is also the year that Roald Dahl published his story "The Great Automatic Grammatizator," in which an engineer named Adolph Knipe convinces his boss, Mr. Bohlen, that they could make a killing by using a computer to write cheap, shitty stories:

> "Nowadays, Mr Bohlen, the hand-made article hasn't a hope. It can't possibly compete with mass-production, especially in this country—you know that. Carpets . . . chairs . . . shoes . . . bricks . . . crockery . . . anything you like to mention—they're all made by machinery now. The quality may be inferior, but that doesn't matter. It's the cost of production that counts. And stories—well—they're just another

product, like carpets and chairs, and no one cares how you produce them so long as you deliver the goods. We'll sell them wholesale, Mr Bohlen! We'll undercut every writer in the country! We'll corner the market!"

Knipe builds the machine and it's like they're printing money. Dahl concludes:

> This last year—the first full year of the machine's operation—it was estimated that at least one half of all the novels and stories published in the English language were produced by Adolph Knipe upon the Great Automatic Grammatizator.
>
> Does this surprise you?
>
> I doubt it.
>
> And worse is yet to come.

It came.

Like artificial intelligence itself, A.I. slop is an artifact of the Cold War. *The U.S. sought to defeat the spread of Communism*, a stingingly grumpy T-shirt might read, *and all we got was the death of books, bookstores, newspapers, and authors.*

Experiments like Strachey's were part of an explosion of postwar research on the relationship between mathematics and language, expressions of a broader fascination with the automation of knowledge, which crossed disciplines and suffused the culture. Among the many unknowns of the Cold War was the extent to which the world was random or ordered. Could the Soviet Union's next move be predicted, or not? "Artificial intelligence" emerged from "intelligence" in the sense of espionage, as computers were deployed to do signal processing—the search for patterns in radio broadcasts and in printed texts like newspapers. Teaching a machine to read becomes useful when you're spying on a twentieth-century enemy. That it could learn to write was a bonus that contributed to a revolution in linguistics and in poetics, too.

In the nineteen-fifties, the fields of computer science and artificial intelligence—both terms were coined that decade—were increasingly concerned with the simulation of human intelligence and with the translation of human (or "natural") language. Linguists were turning language into codes, too. In "Syntactic Structures," published in 1957, the year a science-fiction magazine cover pictured a robot reading a book, Noam Chomsky illustrated the separability of syntax from meaning with the sentence "Colorless green ideas sleep furiously," the kind of thing that might have been written either by the Auto-Beatnik or, to be fair, by an actual Beatnik. Circa 1959, William S. Burroughs started experimenting with writing poetry by cutting up pieces of prose and pasting them together on a page, as in a poem made of newspaper stories about 1) the polio virus and 2) a performance at the Met:

> The girls eat morning
> dying peoples to a white bone monkey
> in the Winter sun
> touching tree of the house. $$$$

At the same time, and using a rather similar method, the German mathematician Theo Lutz created poetry on a Zuse Z22 computer by writing a program that drew from, or cut up and pasted together, random words from Franz Kafka's "The Castle":

> EVERY STRANGER IS FAR. A DAY IS LATE.
> EVERY HOUSE IS DARK. AN EYE IS DEEP.

The Zuse Z22's poetry was reviewed in the *T.L.S.*, where it elicited the opinion that the elimination of meaning was hardly impressive: "What really matters is to eliminate sense."

By the early nineteen-sixties, there was enough of this kind of thing going around that it caused both a panic and understandable excitement. "THE MACHINES ARE TAKING OVER: COMPUTERS OUTDO MAN AT HIS WORK NOW—AND SOON MAY OUTTHINK HIM," a headline in *Life* warned in 1961. What was billed as "the first book of free verse written by an electronic computer" was published in Montreal in 1964, and was credited to "the author, an electronic computer, the LGP-30, which composed the automatic sentences in this collection." Those sentences included this one: "*La pomme ajuste le monde, mais la pluie s'embellit pour les raisins.*" ("The apple shapes the world, but the rain enhances the grapes.") Was it art? Nah, but it was interesting.

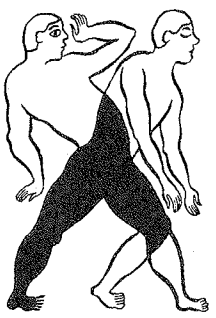

In 1962, the German philosopher and semiotician Max Bense, who had supervised Lutz's work, attempted to draw a distinction between natural and artificial poetry. Natural poetry, Bense wrote, "has as its prerequisite . . . a personal poetic consciousness," whereas, in artificial poetry, there is "no personal poetic consciousness with its experiences, adventures, feelings, memories, thoughts, imaginative conceptions, etc., that is, no pre-existent world, and in which writing is no longer an ontological continuation through which the world aspect of the words could be related to a self." An artificial poem is a poem without a poet.

It's no accident that Bense wrote about artificial poetry, not artificial prose. In Bertram and Montfort's book, the section on poetry is also by far the longest one. Machine-generated text could be baffling: random, and unexpected. Maybe a computer was a new tool for understanding poetry. "You will say that to use a computer to write poetry is like using a crane instead of a pen to write a letter," the British philosopher and computational linguist Margaret Masterman admitted in 1964, but, with the computer, she argued, "we can at last study the complexity of poetic pattern."

Masterman, who studied philosophy and language with Wittgenstein, was, in 1956, the founder and director of the Cambridge Language Research Unit. (Earlier, she'd written novels.) She was a pioneer in machine translation, and her early work established the basic methods of information retrieval. She believed computers could come to understand meaning, and to generate it. She also tried to make that art, producing, with her colleague Robin McKinnon-Wood, "computerized Japanese haiku," like this one: "all white in the buds / I flash snow peaks in the spring / bang the sun has fogged."

Few writers were as enthusiastic about robot writing as Italo Calvino. In his 1967 lecture "Cybernetics and Ghosts," Calvino complained that "the use so far made of machines by the literary avant-garde is still too human," and predicted that a "true literature machine" would someday emerge, one that rejects rules and forms and "itself feels the need to produce disorder." Unfortunately, or maybe fortunately, Calvino's dream of truly original machine literature has yet to be realized: so far, the machine hasn't, in Calvino's

formulation, felt the need to produce disorder, which is to say, literature. Instead, text produced by large language models, however remarkable, sophisticated, and even occasionally wondrous, is derivative, average, predictable. It is language without a mind. But is that even language?

In 1982, in an article called "Against Theory," the literary scholars Steven Knapp and Walter Benn Michaels proposed a thought experiment to show how hard it is "to imagine a case of intentionless meaning." Suppose you're on a beach and discover, written into the sand, this message:

> A slumber did my spirit seal;
> I had no human fears:
> She seemed a thing that could not feel
> The touch of earthly years.

If you didn't recognize the verse as Wordsworth's, you might not worry about its author or its author's intention. You'd just recognize it as writing and try to understand its meaning. But what if, while you were staring at those lines in the sand, a wave came and washed them away, and, when the wave ebbed, it left in its wake another stanza?

> No motion has she now, no force;
> She neither hears nor sees;
> Rolled round in earth's diurnal course,
> With rocks, and stones, and trees.

Well, then you'd have to wonder: Who wrote this, and why, and how? Wordsworth's ghost? The sea itself? God? For Knapp and Michaels, meaning without intention does not exist: "What a text means and what its author intends it to mean are identical." An author without an intention, they argued, is not an author.

Maybe that's a good definition of slop. *Quick bulldozer boost (spam-free!) your toy-truck kid spark + 28-book marathon inspire; let's swap suspense secrets 15 mins this week?*

The end of the Cold War very nearly coincided with the opening of the internet to the public. In the decades since, theorizing of the relationship between natural and artificial literature spawned a whole new academic field, generally within English departments. Courses in what might be described as robot lit are now being offered at universities that include Duke, Columbia, and Harvard. The literary critic Avery Slater argues that computer scientists, military labs, and corporations participated with poets in the Cold War-era creation of what she calls "post-automation poetics," a sensibility that brought together both an artistic vision and an engineering scheme. What was exciting about artificial poetry was that it had no author, no context, no history. It was nothing but form. It therefore had—has—a lot to teach the world about both language and art. A new theory of A.I. slop, however, has yet to emerge, nor a real answer to Wordsworth on the beach. *I had no human fears.*

Long after the fall of the Berlin Wall and the collapse of the Soviet Union, the development of natural-language processing continued in universities, defense labs, and corporate R. & D. departments. Literary experiments with computer-generated text borrowed from that research's developing tools, such as topic spotting. This led to some wacky writing. After the founding of National Novel Writing Month, computer-generated-text devotees founded National Novel Generation Month, in 2013. Leonard Richardson's "Alice's Adventures in the Whale" is a retelling of "Alice's Adventures in Wonderland" with all its dialogue replaced with dialogue from "Moby-Dick": " 'Can't sell his head?—What sort of a bamboozingly story is this you are telling me?' thought Alice." One important tool in early natural-language-processing work was ranking the frequency of word sequences. Using that technique, along with early summarizing tools, the Canadian writer Ryan Stearne cut "The Old Man and the Sea" down to a two-thousand-word short story called "Old Sea."

About a decade and a half ago, the Auto-Beatnik tradition reëmerged on social media, in seemingly automated accounts like @Horse_ebooks, a viral sensation, with posts like: "(using fingers to indicate triangular shape) SMELL SMELL SMELL GOOD NEW NEW NEW slice drink MATCH SPARKLER (thrown in air) STARS STARS STARS," and, most memorably, "Dear Reader, / You are reading." In 2013, *The Atlantic* dubbed the account's output "the Most Successful Piece of Cyber Fiction, Ever." Disappointingly, it was soon revealed that @Horse_ebooks wasn't an automated account but was instead put together by two guys, and I'm not even sure why.

There's nothing wrong with nonsense. But it's not always poetry. And mistaking one for the other is another legacy of how the Cold War foreshortened the humanistic possibilities of the intellectual revolution of the past eighty years—a revolution that has, miraculously, allowed people to communicate with machines using human languages. Shouldn't this be one of the most exciting times in history to be studying language, literature, and literary theory? In "Language Machines," Leif Weatherby, N.Y.U.'s director of digital humanities, points out that, in the years since the Cold War, "the humanities lost language" to cognitive science and computer science. Given that machines can generate language without recourse to reason—he argues that the two things have been radically decoupled—what's needed now is "a theory of meaning in the absence of intelligence." Language no longer distinguishes humans, Weatherby says, dismissing the contention, made by Chomsky and others after the release of ChatGPT, that L.L.M.s "differ profoundly from how humans reason and use language." Weatherby calls this, in a curious choice of metaphor, "remainder humanism": "a humanism without a theory or doctrine of what is human, in which humanity is remaindered, like a book past salability."

If that's what it means, now, to be human—to cling to the idea of a relationship between language and reason—I don't think I mind being a book, even a remaindered one, shelved in the dark downstairs of a bookstore and priced cheap. Is the alternative really so enticing? This winter, the most popular series on TikTok was reportedly "Fruit Love Island," an entirely A.I.-generated version of "Love Island" featuring talking fruit. "Welcome to Fruit Love Island, where eight single fruits are about to flirt, fight, and trust," it goes. DEAR HONEY-DEW YOU ARE MY GREATEST WHISKERS MY UTTER MOONBEAM. . . . YOURS BEAUTIFULLY MANCHESTER UNIVERSITY COMPUTER.

Something big is happening, something fascinating: we can talk to machines. "We do not have language yet for this twist in our plot," Weatherby writes. The contention is that something big is happening to us, that someone else, something else, is writing the plot. But shouldn't we be writing it? Because, so far, that plot is slop. ♦

HISTORICAL / GETTY