

OPINION

The automation of writing is almost here. But what will happen to us if we cede our written language to AI?

MICHAEL HARRIS

SPECIAL TO THE GLOBE AND MAIL

PUBLISHED NOVEMBER 22, 2024

FOR SUBSCRIBERS



A robot writes out the Bible using precise calligraphic script in Robotlab's 'bios' art installation.

WIKIMEDIA COMMONS

 34 COMMENTS

 SHARE

 BOOKMARK

 GIVE THIS ARTICLE

Michael Harris is the author of several books, including Solitude: A Singular Life in a Crowded World and The End of Absence: Reclaiming What We've Lost in a World of Constant Connection.

In one or two years I'll have the chance to never write again. You, too, will have this opportunity. Generative AI will deliver floods of human-free text and even a human-free framework for all communication. We are strolling onto a grand (and poisonous) landscape where our own thoughts and intentions are leached from our words.

One [study](#) finds that, already, the majority of online content is generated by AIs and experts predict that 90 per cent of all content [will be AI-generated](#) in a year or so. As each new AI model consumes and regurgitates this content in its turn, tomorrow's text could steadily degrade, like a memory of a memory of a memory. [New research](#) from Ilia Shumailov at Google DeepMind finds that AI models that are trained on the output of earlier AIs will eventually devolve into a "model collapse" that produces "irreversible defects." The New York Times recently [reported](#), for example, that legible numbers melt into meaningless cyphers when AI models regurgitate them a few dozen times. In fact, our brilliant AI systems may soon "choke on their own hallucinatory fumes," [says](#) cognitive scientist Gary Marcus.

Our written language has always been blurred by vague usage, clichés and careless generalities. But the advent of AI-generated text brings a far riskier degradation. If AI dominates our written culture, and if floods of AI writing become the training material for future models, then the words that guide and shape our text-gorged lives will creep further and further from our own. We will bandy ready-made phrases back and forth – prefab words that are shipped and passively consumed but lacking in the grist of human intention. To what end? George Orwell warned us long ago that the effect of ready-made phrases is always the same: They anesthetize the brain.

We ingest this anesthetic each time we read ChatGPT's writing. There's a heartless, droning circularity to its prose that seems to portend a larger somnambulism. There is the shape of meaning and yet nothing solid. There's a synthetic flavour of intention laced over a sociopathic disregard for truth. Yet we're quickly growing used to that grey style. And, since ease will always be prioritized over effort, the automation of our writing will, one day soon, be so commonplace that idiosyncratic (i.e. human) styles will feel odd and grating.

But what exactly will happen to us if we cede our written language to AIs that do not – cannot – believe what they write? What happens to our daily communications and expressions when the elements of genuine, personal style are replaced by those memories of memories of memories? What new banality is born?

There have been three great crises in the history of authorship. The first came with the dawn of writing itself, and our divorce from the age of oral storytelling. The second came with the dawn of the printing press, when books (and authors) became marketable. The third crisis, which has only just torn into our

collective experience, is the dawn of AI. If the first two crises created individual authors, this third crisis dissolves them. Shortly we'll have no idea whether a comment, e-mail, report or novel was created by a single human mind or whether it's a mulchy regurgitation of all the comments, e-mails, reports and novels that came before.

This change is massive and we've only begun to grasp its effect. Every change in our communication technology realigns our system of knowledge, changes our knowledge monopolies – often in unexpected ways.

When papyrus became cheap in the 7th century, for example, there was a flowering of Greek lyric poetry. When books became a mainstay in 16th-century Europe, grand architecture held less meaning. It's difficult to anticipate these changes but we do know that AI's realignment of knowledge involves a draining away of human intention and a concurrent inflation of available text.

There's a name, of course, for an excess of verbiage without any real intention behind it: We call that bullshit.

Bullshit may, in fact, be the defining characteristic of the new dominant style. Auto-generated text is the antithesis of sincerity, after all, because the AI cannot care about what it's sharing. When the philosopher Harry G. Frankfurt wrote his treatise *On Bullshit*, he noted that such an "indifference to how things really are [is] the essence of bullshit." And yes, when reading AI-generated text we can recognize its cool indifference; somewhere in our language-obsessed brains a warning flares.

Only for a moment, though. Only for this brief moment in time when the weirdness of AI text is still apparent. Very soon the aloof tone will be suffused into our mental atmosphere. It will be normal. Style will be flattened, collectivized, and from the various histories of human literature will emerge a single mush. Our children, of course, will not know the difference. In fact, they'll think the AI's dull bias is also their own.

Do I catastrophize? Well, bullshit is dangerous. Words are vehicles for thought; to flatten the language is to flatten our minds.

Meanwhile: Our heroic tech plutocrats say they've freed us from the anxiety of language, its infinite unfair vagaries. You are free, they promise, from the

stress of composition; free from the risk of miscommunication, the red-faced shame of being misunderstood; free from analysis, too, and from scouring overlong documents; free from searching for the word, the phrase, the exact idea; free from cobbling together the sentiments that now appear like magic on your screen.

Does this cognitive offloading actually feel like freedom, though? Researchers find, unsurprisingly, that those who automate mental tasks tend to [turn passive](#) and their engagement withers. Passing around prefabricated blocks of text is efficient, to be sure, but it also makes a mockery of the word “communication” and could exacerbate our [loneliness epidemic](#). Real communication (which is to say, real connection) requires labour; it requires us to think, to search inside of ourselves, to reveal. The author Nicholas Carr warns that labour-saving devices don’t simply provide a shortcut for our jobs.

They “alter the character of the entire task, including the roles, attitudes, and skills of the people who take part in it.” In other words, automation is always a bargain. You are given ease, for example, but loneliness is attached. You are given comfort but aimlessness comes too.

It is, of course, nearly impossible to measure the cost. The erosion of autonomy and engagement that AI requires as it absorbs our written culture will be as mysterious as it is profound.

But we can begin to measure the scale of that cost by recalling the scale of written culture itself. For better or worse, writing has been our civilization’s driver and definer. We codify, we document, we classify and record. Writing superpowers the human brain’s natural abilities, magnifying knowledge and democratizing access. Each revelation and misunderstanding, each expression of love and hate, each moment of learning and forgetting, all of it is achieved through our words. The linguist R.F.W. Smith put it more succinctly when he said that we live our lives in “a house that language built.”

And so, to toss our language into the echo chamber of AI is to abdicate part of our larger selves as well. If the hallmark of AI-generated text is that it regurgitates without intention to the point of being a bullshit machine, then we must accept that lives run and defined by AI-generated text will grow vague and bullshit-y too.

I don’t think we want to be simply labelled and shuffled into the morass of some large language model. I think every one of us longs to be properly named

and intimately described. Yet we're building a future where we aren't even the narrators of our own story.

Maybe that demotion is the point.

When Orwell conceived of Newspeak – the “double-plus-good” language in Nineteen Eighty-Four – the idea was to streamline and condense all words in order to streamline and condense human thought. Orwell elaborated on the strategy in his essay “[Politics and the English Language](#),” where he explains that dulling language with euphemism and clichés and vagueness (those hallmarks of generative AI) is the easiest way to tamp down political resistance. The vagueness produces a smokescreen against action; it even allows a “defense of the indefensible.”

This is the true danger posed by successive rounds of AI-generated language – that our ideas become so ready-made and so passive that we merely prompt prose and never compose it ourselves. Then we won't have the skill or wherewithal to name tomorrow's horrors, let alone argue against them.

When an autocrat murders his people, when the rights of workers are gutted, when facts themselves begin to blur, will we know how to call the problem out? If lazy generations of upchuck language have replaced the hard work of real writing, then perhaps not. Systems of communication, after all, are not just tools for the noble expression of human thought; they can also be tools for thought's repression.

<https://www.theglobeandmail.com/opinion/article-the-automation-of-writing-is-almost-here-but-what-will-happen-to-us-if/>

We know that if there is money to be made in it, and the ecological and societal costs can be confined ostensibly to “somewhere else” or obfuscated, a dangerous technological innovation will be swiftly put in place. Look at all the server farms being set up to serve AI, blockchain, and chatbot purposes; we know that the humungous electrical energy required to power them, however derived, is not good for the planet. Employers, driven to satisfy investors (including themselves) in a competitive marketplace, have always preferred machines to people; the latest form, a chatbot or robot, is much easier to manage than quirky people prone to problems. As long as enough people keep reaping enough from this method of production to purchase the products and services robots and code pump out, the further dehumanization of industry will continue and deepen. Even those disadvantaged or displaced will accede to it because it will let them stretch their lessened meagre earnings farther. TJB