Machine Writing, Learning, and the Disappearance of the Pen

Andrew D. Digh
Mercer University, Macon, Georgia

Kevin Cummings
Mercer University, Macon, Georgia

ABSTRACT

Writing is an evolving technology. When Walter Ong famously explored the differences between oral and written cultures, he predicted electronic technologies are inaugurating a new era of secondary orality. Drawing from Ong, Yohei Igarashi takes up the question of how machine writing reconfigures written communication. Students today receive editorial assistance in their work through programs such as predictive text and spell check. The increasing sophistication of computer programs that can edit and author (or co-author) written works poses numerous challenges to the university. Our essay explores tectonic shifts taking place with new forms of machine writing and artificial intelligence. Machines are colonizing educational spaces and already provide substantial assistance to students. Teaching students to acquire the habits of critical reflection essential to being critical thinkers requires faculty to become more multimodal in their pedagogy and develop a repertoire of assessment tools that go beyond traditional summative assessments of student writing.

1. Introduction

We teach students the craft of writing, so they can discover and unravel the world around them and wander into imaginary places. Because words create worlds, they have power and magic. Today, computer programs usurp the place once held solely by human authors by becoming inventors of stories and arguments. In this essay, we examine how machine writing challenges our system of modern education. Students now co-author work along with machines capable of far more than just spell-check and predictive text. Machine writing programs are creative partners. This creates a unique challenge for faculty who aim to teach students writing as a way of cultivating habits of critical reflection. Our interests in this project nest within three disciplines: Communication, Composition, and Computer Science. Each is radically transforming the paradigm shifts that machine writing has inaugurated.

Yohei Igarashi (2021) maintains that the work of Walter Ong (1912–2003) is indispensable to understanding what machine writing will mean. Ong famously argued that writing restructures consciousness. More than just a technology that advances human knowledge, Ong claims that writing transforms how we make sense of the world. Igarashi reminds us how Ong understood preliterate cultures (Ong 1982). Their reliance on memory was coded directly into the language they used. This functioned as a precursor to literate cultures, which Igarashi tells us became freed from needing to have memory aids inscribed into words, because they had a reservoir of memory in writing—a technology that is also an archive of human thought. The exponential expansion of capacities for information storage allowed for in writing generated new modalities of creativity. Igarashi anticipates machine writing may similarly produce parallel opportunities for new ingenuity.

In the context of higher education, concerns about plagiarism and a generation of students who may not understand the basics of writing cast a large shadow over Igarashi's optimism about machine writing. Many faculty rely on programs such as Turnitin to catch academic dishonesty. As machine

writing produces work of high quality that is original in form, students have fewer temptations to cut-and-paste paragraphs from the work of others. In addition, they need no longer fear that their essay will already exist somewhere in the system. Machine writing programs help produce original work that current anti-cheating programs cannot trace. To confront the challenges of machine writing, we think most faculty will need a more adaptive and multimodal pedagogy, one reliant more on active learning and formative assessments of writing and less on singular, summative assessments of final written products.

Our argument begins with an introduction into current iterations of machine writing in order to illustrate what these new engines of innovation can accomplish. This then provides a fertile space for returning to the work of Walter Ong and rethinking how writing curates and transforms thought. Having set the stage for a cross-pollination between machine writing and systems of meaning-making, we will turn to education and take up the pedagogical obstacles that faculty face in the new technological milieu. Finally, we set out some of what we believe will be preliminary best practices for using the new machine writing tools.

2. Exploring a New Technology

Machine-generated writing technologies have now been used successfully in a plethora of disciplines, including the arts, the humanities, and the social sciences over the last five years. They all utilize software that *learns* to produce *human-level* text by analyzing enormous amounts of text on the internet, processing what letters and words tend to follow each other. This software—known as Generative Pre-trained Transformer (GPT)—was first developed by a team of scientists at OpenAI, an artificial intelligence research lab founded by Elon Musk and Sam Altman, and based in San Francisco, CA. It was originally created in order to "help humanity" (Lomas 2022: para. 5). GPT is largely based on neural networks. A neural network is a method of artificial intelligence that models the structure of the human brain. It is a machine learning technique which features interconnected nodes in layers similar to those found in the brain.

The third iteration of the original GPT, known as GPT-3, can process natural language tasks that range from language translation to generating fiction or poetry to predictive text to answering any type of question. It tries to match the context or pattern in the text you provide with information found in similar or integrated writings. GPT allows people with little knowledge of programming to get quick access to machine learning software. As you can see in Figure 1, there have been tremendous strides in natural language processing input in recent years (Das 2020). GPT is the largest language model ever built, encompassing 200 billion machine learning parameters (Radford 2019).

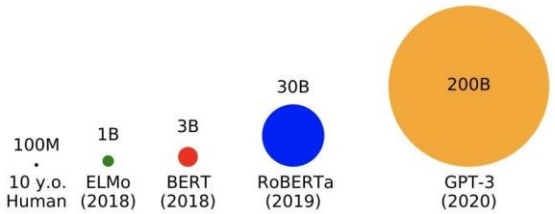

Figure 1: Comparison of human and model linguistic input (# of word tokens)

One of the disadvantages of GPT-3 is that it does not have the so-called "brain power" to understand what exactly it has written. For example, if I type the sentence: "Andy and Kevin are co-authoring an article together. They have really enjoyed writing it." and then start to type a second one, GPT-3 can help me finish writing it, but in doing so, GPT-3 struggles to decipher the meaning of the words "they" and "it". This is often connected to Moravec's Paradox, which says that the simplest tasks for humans are often the hardest for machines and vice versa. That is, tasks many humans find challenging like solving a calculus problem or playing chess are incredibly easy to program in AI. However, consider these answers given by GPT-3 to similar, but different queries that are easily handled properly by humans:

Q: How many eyes does a giraffe have?
A: A giraffe has two eyes.

Q: How many eyes does my foot have?
A: Your foot has two eyes.

Q: How many eyes does a spider have?
A: A spider has eight eyes.

Q: How many eyes does the sun have?
A: The sun has one eye.

These queries show a lack of common sense on the part of the program, as well as a failure to pass The Test—a classic AI test of intelligence (Lacker 2020). Some, however, might say that this shows that even a machine can have a sense of humor!

In the natural language processing diagram on the left-hand-side of Figure 2 below, you see input—fed through a GPT-3 model—that is processed one word or "token" at a time. Fine tuning via further training continues to improve this type of machine learning over time, so that GPT programs can produce better output, furthering their apparent "understanding" in the future. The GPT training is actually using unsupervised learning implemented by what is called a *deep learning* network like the one you see pictured on the right-hand-side of Figure 2. Deep learning can be thought of as machine learning with additional power.

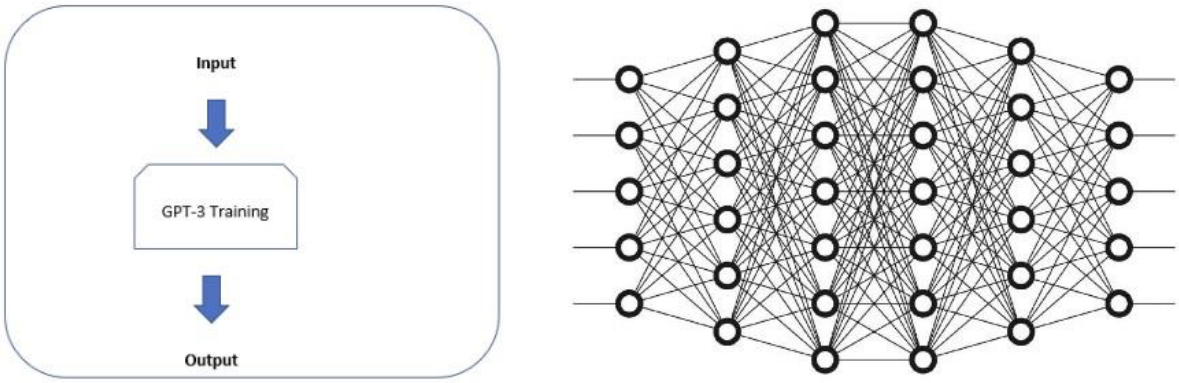

Figure 2: GPT-3 takes in input one word at a time and feeds the words into a training module. This module continually improves itself using a vast network of neurons that talk to each other based on information already learned. These networks are based on information processing modules in real biological systems.

In traditional machine learning, you have supervised learning where all the training happens on pre-labeled data like a dog, house, or car. With deep learning, we have multiple layers of artificial neurons that feed information into other neurons to progressively extract additional information. There is a ton of data out on the internet that is not yet labeled. It is the goal of GPT to essentially better label the information on the internet and learn from that data in an unsupervised fashion. If it does so successfully, the pace of the AI can be accelerated and the text produced can become more human-like because it has (to relate it to human terms) "read more books". That is, it has more ideas to connect with, extract from, and write about. This correlates well with a famous quote by Virginia Woolf that is on display at the Library of Congress: "Read a thousand books and your words will flow like a river" (Woolf in Hansford 2020: para. 2).

A current constraint of GPT-3 is that it has a limited input size. It can take in 2,048 linguistic tokens, or about 1,500 words. However, that is a significant improvement over earlier iterations of GPT. Also of concern are the data sets on which the GPT models are actually trained. Some have been found to "output hateful language, replicating the kind of writing from which they've learned" (Igarashi 2021: para. 4). It is important to remember that GPT-3 is "just making a guess based on statistical patterns in language, and that may or may not have any correlations to the world as humans understand it" (Dzieza 2022: para. 12). These limitations remind us that "we need to take it in context: it doesn't work like the brain, it doesn't learn like a child, it doesn't understand language, it doesn't align with human values, and it can't be trusted with mission-critical tasks." (Marcus 2022: para. 35).

Similar concerns have been raised about DALL-E, a machine learning technology developed by OpenAI that can be used to create artwork from text. Its most recent model was trained on 650 million images scraped from the internet paired with text captions (Johnson 2022: para. 2). Many of the artistic images "lean toward generating images of white men, overly sexualize images of women, and reinforce racial stereotypes" (ibid: para. 3). Some artists are getting understandably nervous about some of the incredible images produced by another similar AI program called Midjourney. The issue can be seen in Jason Allen's award-winning image "Theatre D'opera Spatial". People enter a string of words into Midjourney and almost instantly, they receive an image in response (Roose 2022: para. 9) The aforementioned image generators work their magic using deep learning algorithms that are applied to visually pleasing images obtained from the web. They are essentially producing a mishmash of art pieced together from works already made by real life artists.

AI also seems to be on a course to truly change how we view written media. One example of this can be seen in AI-written news stories. Many thought computers "would never be capable of doing the more creative, human parts of our jobs: generating story ideas, gathering quotes from reluctant sources, or explaining complex ideas in an accessible way" (Roose 2021: 63). However, robo-journalists can take structured data and "turn it into a full-fledged news story in milliseconds" (ibid:63). Some journalists have even been replaced by a natural language generation app, Wordsmith, which can produce "300 million news stories in a single year—more than every other media outlet in the world combined" (Roose 2014: para. 12). In sports, so-called "robot journalists" can now easily take a box score and produce a game recap. However, AI struggles to show a human touch, especially "when using metaphors, humor, and poetry" (Ombelet 2016: 732).

3. Thinking About Writing

As machine writing programs become more and more capable, the dissemination of technology-enhanced work will become common. Writing historically is a transformational technology and revisiting the history of how writing has influenced knowledge production can help us consider where the new paradigm shift will take us. If we consider the biography of writing, it is clear that this technology has already altered human destiny. Writing is a unique human invention that allows our species a means to store knowledge. This opened up possibilities for human activity. Ong, in his work *Orality and Literacy: The Technologizing of the Word* explains, "Technologies are not mere exterior aids but also interior transformations of consciousness, and never more than when they affect the word" (1982: 82). Through writing, new kinds of thought processing advanced as utterances are reformulated into what Ong identifies as "exquisite structures" (ibid.: 84). For students of semiotics, the new pathways for making and recording meanings expanded exponentially. Moreover, importantly, the functions of words transform as they transfer from the channels of speech and voice into the channel of writing.

Igarashi is especially interested in how writing records knowledge and the refrain he returns to repeatedly is the idea of storage. Writing stores knowledge. It is a container for thought. In understanding the influence of writing as a tool, Igarashi avoids the pessimism that condemns machine writing to a role supplanting human creativity and labor. Instead, he thinks composition is playfully expanded by new technology. Igarashi contends humans can productively partner with machine writing programs and the result allows for new innovations and productive creative work. This is an idea we return to later.

To understand what the paradigm shift will mean for writers, it is perhaps useful to review how machine learning has already influenced other activities. The study of mathematics at advanced levels routinely relies on machines from calculators to computer programs. In the game of chess, a program on a phone provides a level of play that surpasses the greatest grand masters. In Theater, set design is dramatically enhanced by sophisticated artificial intelligence programs. Jason Jamerson, in an interview with Michael Schweikardt, explains "Up until now, we'd read the script, look at a bunch of research, and venture alone along a dark road to imagine a synthesis of our needs and passions for a design. Then one day we'd turn up at a meeting with sketches, ground plans, etc. Now, for the first time, thanks to AI networks, we can bring something with us on that road where our ideas melt together. I think AI is a lantern" (Jameerson in Schweikardt 2022: para. 5).

4. Toward a Multimodal Pedagogy

When we ask students to become proficient in writing, what we really want is for our students to acquire the ability to make their thoughts intelligible and accessible to others through the medium of the written word. It is through the sharing of knowledge that our species developed. Schooling produces certain kinds of individuals. Because education helps form us as citizen subjects, humans are experiencing a renaissance of wonders. Writing is a fundamental process for nurturing habits of critical reflection. We want our students to learn the craft because it advances our capacities to think collectively together to solve problems great and small. Universities are havens for imagination and creative work and offer students amazing opportunities to learn and advance. The possibility that artificial intelligence creates risks for the rigorous intellectual labor students engage in as part of learning is certainly one that should have all of us worried. If students produce AI assisted essays that require little labor or effort on their part, that short circuits their learning process and leaves them unprepared for the challenges to come. We believe the emphasis in higher education among faculty needs to shift from a focus on summative assessments of final essays, toward a pedagogy that deploys scaffolded learning opportunities, uses formative assessments, and requires students to demonstrate learning of appropriate and inappropriate uses of machine writing technologies.

Lev Semonovoch Vygotsky provides a useful entry point for educators to consider as we reckon with how to navigate a future where students can utilize technology to fabricate their arguments. He writes, "The relationship of thought to word is a vital process that involves the birth of thought in the word. Deprived of thought, the word is dead" (Vygotsky, 1987: 284). When we judge the capacities of a student by the words in their essays, we risk using remnants of thoughts that are more cybernetic than human. We judge a script that no longer represents their individual creative process. To capture and understand the vital natal process Vygotsky describes, we need to find new ways for students to articulate and share their ideas. We cannot simply collect papers at the end of a semester and expect to be able to make an informed judgment of their labor and work.

Our proposal is simple. First, university education should adapt away from the exclusive use of essays students write in isolation on computers as the objects of our assessment. Writing has intrinsic value, and we think essay writing in particular has served our learning communities well for quite some time. However, the new milieu requires that we have students go back to using pen and paper if we truly want our young authors to originate their own work. Second, we should be careful not to assume that technology is ruining education and try for a ban. Math instruction today exceeds the work done in the past because of machines such as calculators and computers. With proper use, writing proficiency will also become easier. We all already benefit from basic programs. More advanced programs will no doubt pose challenges for us as educators, but there is no world where the genie is placed back into the bottle. Instead, we need to require more hands-on learning that encourages students to find good uses for machine assistance.

5. Co-Authorship

David Ferrucci, the IBM researcher that helped lead the development of Watson, the AI system that played and won Jeopardy, believes AI should become "a trusted thought partner" for us in the future (Ferrucci 2017:00:01:08-00:01:50). This is becoming ever closer to a reality thanks to Sudowrite, an AI writing assistant for "unlocking creativity and breaking writer's block" (Robertson 2023: para. 1). It is currently assisting popular fiction writers in meeting tight deadlines where "eager readers expect a new novel every nine weeks" (Dzieza 2022: para. 1).

A screenshot of the Sudowrite interface can be seen (in Figure 3) after it has been given an informal writing prompt related to the novel *Lord of the Flies*. This is a prompt assigned to college students in a first-year, integrative seminar course. Notice Sudowrite's "soothing sunset-colored interface" (Dzieza 2022: para. 6) and ease at assisting the writer in describing, rewriting, and brainstorming. The designers of Sudowrite, Amit Gupta and James Yu, "collected plot twists from short stories" and fed them to GPT-3 along with sentence descriptions of related "smells, sounds, and other senses" (Gupta and Yu in Dzieza 2022: para. 14). Sudowrite can provide many of the same services as a real writing tutor. In addition, it can help students scaffold their writing over time including the gathering of more pertinent sources as well as help in revising their work.

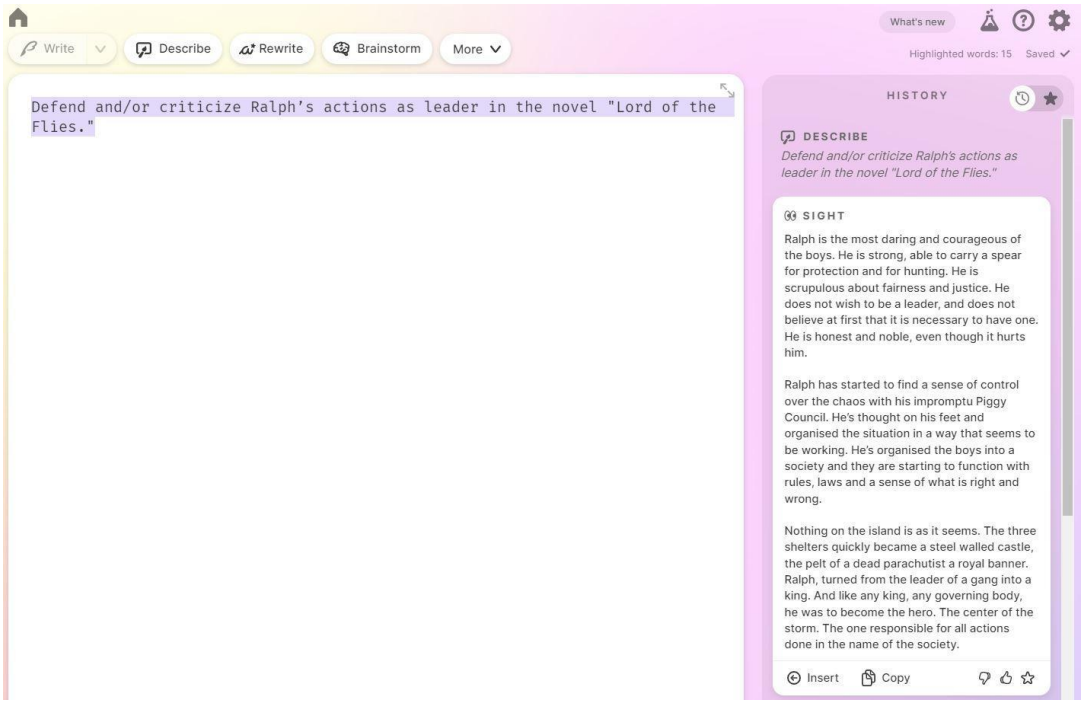

Figure 3: A screenshot example of Sudowrite in action with a Sorbet-themed writing background, a high level of creativity, and responses of approximately 250 words. At the top, it includes options to describe, rewrite, and brainstorm. Notice the clear, concise writing in the SIGHT box on the right. This response was one of two different choices generated in approximately one second to the prompt on the left.

In 2020, GPT-2 was successfully used to create the popular, humorous story *Harry Potter and the Little Prince*. The book was written by Botnick, a machine entertainment company whose products are at the intersection of AI and comedy. GPT-2 was the perfect model for what Botnick wanted. It already knew language fundamentals, so they were able to just instill further training on a select data set that they wanted it to imitate. In the case of *Harry Potter and the Little Prince*, the program was trained on the series of the seven Harry Potter books written by author J.K. Rowling. Although it seemed insignificant at the time, the writing of this AI-generated novel helped plant the seed of how AI could change how we

view written media. When we hear the term "AI-generated novel", many may feel skeptical. In the past, humans were the only entities writing, and the actual process of writing was "fully immersive" (Lepp in Dzieza 2022: para 40). That seems to be changing now. Many writers of fiction now report that, after using Sudowrite, they no longer "wake up thinking about their characters" and their "writing does not feel like their own anymore" (Dzieza 2022: para. 42).

The definition of intelligence often used by psychologists focuses on skills that humans have, such as abstract reasoning, mental representation, and creativity (Ruhl 2020). GPT-3 seems to be meeting this definition by "manipulating higher-order concepts and putting them into new combinations, rather than mimicking patterns of text" (Johnson 2022: para. 43). It displays this budding intelligence when it is given a prompt to compose an essay comparing two completely unrelated things. For example, when I gave Sudowrite the prompt: "Compare the music of Bob Dylan to a computer", it came back with this interesting response, which shows that it can generate new, commonly understood ideas.

> Music doesn't look like anything, of course. But to compare the music of Bob Dylan to a computer, you might think of it like a vast, white room devoid of any interesting features. There are no chairs, no place to sit. There are no windows. The ceiling is high and made of a clear, tinted material, so you can't look up. Every wall is white and blank. In the center of the room, in a large clearing, there's a spotlight shining on a single object: A computer.
>
> Dylan is sitting alone at a keyboard. He's playing a new tune he's been composing. It's like a painting, a collage of images and pieces, songs that seem to blend together. It's like a vague memory of summer camp, the smell of a rotting log, and an image of a girl with a guitar and a pensive smile.
>
> A computer is rectangular and black, a flat screen with a keyboard or touchpad. Dylan's music makes use of chords and the voicing of melody; it is the combination of notes that makes it distinct from others (Sudowrite 2022: Sept. 28).

This high quality, articulate essay makes several correct higher-level connections, including the fact that music does not look like anything, as well as the creation of a simile comparing a Dylan song to a painting. It is truly difficult to know for sure if this was written by a person or a machine. It is quite descriptive and uses sophisticated vocabulary. There are no spelling or grammar errors. If you gave high school students a similar prompt, many would be hard pressed to compose a stronger response. They would most certainly not be able to do so in the single second it took Sudowrite to deliver this upon request. When you see this program writing sentences of this caliber, it is hard not to feel that the machine is thinking in the same way that sentient organisms—like humans—do. Also, it makes you realize that "large language models are going to pose huge challenges for educators trying to prohibit plagiarism—assuming it's still considered plagiarism if a machine writes an essay for you" (Johnson 2022: para. 42).

6. Conclusions and Future Work

In the poem "how we are connected," poet W.E.R. LaFarge writes that language "threads us through and through" (LaFarge in Hess 2008: p. viii). Writing was once a defining quality that made humans distinct from other non-humans; however, this is no longer the case. Today, if you ever find yourself stuck while trying to compose the ending of a poem, essay, op-ed, or book review, you can simply ask Sudowrite to finish it for you. *The New York Times* columnist Kevin Roose did just that recently in "a book review" aptly titled "A Robot Wrote This Book Review". AI-powered automatic machine writers, like Sudowrite, use a GPT deep learning technology that promises to get better and better. The OpenAI team that first developed it is now testing it on supercomputers with exaflops of computing power that continue over longer ranges of time. Future GPTs will make our private digital assistants, such as Alexa and Siri, look like a child's toy. As Roose laments, "Although today's AI systems can be clunky and erratic at times, they are getting better fast, and will soon match or surpass human proficiency in a number of important tasks, solving problems in ways no human would have thought to solve them" (2022: para. 9).

The Chronicle of Higher Education recently published an article entitled "Will AI Kill College Writing?" (Schatten 2022). It forecasted that "it won't be long before GPT-3, and the inevitable copycats, infiltrate the university" (ibid.: para. 8). It also verified that the writing that is output by this technology is original in that it "cannot be detected by anti-plagiarism software" (ibid.: para. 9). A planned future study in the next two years will examine the effects of this software on both university students and professors. This should prove to be very important. Research that also closely monitors the knotty array of ethical issues encountered in the use of this technology will be paramount. The issues surrounding GPT writing include misinformation, author forgery, and bias in word associations.

In November 2022, OpenAI released a new, large language model chatbot to the general public known as "ChatGPT". It will be based on GPT-3.5 and will interact in a conversational way, using a discussion format. It can also argue and write "jokes (some of which are actually funny), working computer code, and college-level essays" (Roose 2022: para. 6). It has been touted as "the best AI chatbot" (Rebelo 2023: para. 4) capable of producing answers that are "cogent, well-reasoned, and clear" (Tufekci 2022: para. 6). A recent comparison between fourth grade writing and ChatGPT output has already fooled many experts on children's writing, including Judy Blume. "None of them could tell every time whether a child or a bot wrote the essay" (Miller 2022: para. 4). Although ChatGPT still has "plenty of room for improvement," (Grant 2022: para. 3) it is easy to see why it recently caused management at Google to issue a Code Red. In their minds, "ChatGPT represents the arrival of an enormous technological change that could upend its search business" (ibid: para. 3).

Does GPT actually put us on a path toward conversational machines and truly sentient systems? This is often referred to as the day that the Technological Singularity or Artificial General Intelligence will be first established. Some have forecasted that this would happen as early as the year 2030, while others predict it by 2045. Many think this will be the moment in time when machines become more intelligent than the humans who created them. It is certainly possible that GPT will become an explosive technology, much like the Greek alphabet was in fostering a golden period of so many great scientific, philosophical, and literary achievements in ancient Greece. Others feel it may be opening a Pandora's Box that could lead us toward a technological dystopia.

No matter what happens, many writers have already started adjusting their own approaches to machine writing and plan to keep GPT technologies on a short leash. For example, writer Jennifer Lepp says she still uses Sudowrite but that she first "pastes everything she's written so far into it, leaves a sentence half-finished, and only then lets it write" (Lepp in Dzieza 2022: para. 59). To her, it is a type of research assistant that can help enrich her writing. Having this tool as a thought partner to complement our work, as well as using it as a tutor, are among the most exciting future promises to consider. At the very least, GPT is sure to become a new tool for "building all sorts of new technologies and new products that could transform the human experience" (Metz 2020: para. 50).

Human machine interactions inaugurate a new milieu that challenges educators. The study of systems of signification provides us with a unique outlook on this moment. Umberto Eco (1986: 137) suggests a key question to guide our thought by asking: "Is the chemical composition of every communicative act the same?" Searching for an alchemy to interpretate human's messages, Eco would have us understand that meaning exists in more than just the medium. Each member of the audience also brings their own understanding into the cauldron of human experience. As *we* interpret communication, Eco has a message about teaching that we think anticipates and helpfully amplifies our own ideas. In writing about the internet, Eco argues, "In the end, a good teacher always notices when a text has been copied indiscriminately and will sniff out the trick" (2017: 61). We think that lesson is equally relevant today. A heavy dose of good teaching helps make machines into useful tools and good thought-partners. Good teaching is essential to navigating this moment in history.

References

CHALMERS, David.
2020. “GPT-3 and General Intelligence”, *Daily Nous*, 30 July 2020, dailynous.com/2020/07/30/philosophers-gpt-3/#chalmers(accessed on 1 June 2023).

DAS, Sejuti.
2020. “GPT-3 Vs BERT For NLP Tasks”, *Analytics India Magazine*, 11 Sept. 2020, analyticsindiamag.com/gpt-3-vs-bert-for-nlp-tasks (accessed on1 June 2023).

DZIEZA, Josh.
2022. “How Independent Writers Are Turning to AI”, *The Verge*, 20 July 2022, www.theverge.com/c/23194235/ai-fiction-writing-amazon-kindle-sudowrite-jasper (accessed on1 June 2023).

ECO, Umberto.
1983. *Travels in Hyperreality*, trans. William Weaver (San Diego: Harcourt).
2016. *Chronicles of a Liquid Society*, trans. Richard Dixon. (Milano: La Nave Di Teseo).

FERRUCCI, David.
2017. “David Ferrucci - Machines as thought partners”, BigSpeak Speakers Bureau, 7 Aug. 2017, www.youtube.com/watch?v=vsyPZdt6noE (accessed on 5 June 2023).

GOLDING, William.
2011. *Lord of the Files*. (New York: Perigee)

GRANT, M. and C. METZ.
2022. “A New Chat Bot is ‘Code Red’ for Google’s Search Business”, *The New York Times*, 21 Dec. 2022, www.nytimes.com/2022/12/21/technology/ai-chatgpt-google-search.html (accessed on 1 June 2023).

HANSFORD, ANNALISA.
2020. “Why Reading is Important for Your Writing”, *Medium*, 12 July 2020, https://medium.com/inspired-writer/why-reading-is-important-for-your-writing-214ae904ecb9 (accessed on 25 May 2023).

HESS, Elizabeth.
2008. *Nim Chimpsky: The Chimp Who Would Be Human* (New York: Bantam)

IGARESHI, Yohei.
2021. “The cliche writes back”, *Aeon*, 9 Sept. 2021, aeon.co/essays/machine-writing-is-closer-to-literatures-history-than-you-know (accessed on 1 June 23).

JOHNSON, Kari.
2022. “DALL-E 2 Creates Incredible Images–and Biased Ones You Don’t See”, *Wired Magazine*, 5 May 2022, www.wired.com/story/dall-e-2-ai-text-image-bias-social-media (accessed on 1 June 2023).

JOHNSON, Steven.
2022. “The Writing on the Wall”, *The New York Times Magazine*, 17 Apr. 2022, 28–35, 44–49.

LACKER, Kevin.
2020. “Giving GPT-3 a Turing Test.”, in author's blog, 6 July 2020, lacker.io/ai/2020/07/06/giving-gpt-3-a-turing-test.html (accessed on 1 June 2023).

LOMAS, Natasha.
2022. “ChatGPT shrugged”, *TechCrunch*, 5 Dec. 2022, www.techcrunch.com/2022/12/05/chatgpt-shrugged/ (accessed on 25 May 2023).

MARCUS, Gary.
2022. “The New Science of Alt Intelligence”, *Substack.com*, garymarcus.substack.com/p/the-new-science-of-alt-intelligence (accessed on 1 June 2023).

METZ, Cade.
2022. “The New Chatbots Could Change the World. Can You Trust Them?”, *The New York Times*, 11 Dec. 2022, www.nytimes.com/2022/12/10/technology/ai-chat-bot-chatgpt.html (accessed on 6 June 2023).
2020. “Meet GPT-3. It Has Learned to Code (and Blog and Argue)”, *The New York Times*, 24 Nov. 2020, www.nytimes.com/2020/11/24/science/artificial-intelligence-ai-gpt3.html (accessed on 1 June 2023).

MILLER, C., A. BUCHANAN, L. PLAYFORD, and A. KROLIK.
2022. “Did a Fourth Grader Write This? Or the New Chatbot?”, *The New York Times*, 26 Dec. 2022, www.nytimes.com/interactive/2022/12/26/upshot/chatgpt-child-essays.html (accessed on 1 June 2023).

NICHOLS, Angela (Ed.).
2020. “Harry Potter and the Little Prince”, GPT-2 generated by *Botnik.org*, 14 Apr. 2020, botnik.org/harry-potter-and-the-little-prince/ (accessed on 21 June 2023).

OMBELET, Pieter-Jan, Aleksandra KUCZERAWY, and Peggy VALCKE.
2016. “Employing Robot Journalists”, Proceedings of the 25th International Conference on World Wide Web, Montreal, April 2016, in *Companion Volume*, ed. Jacqueline Bourdeau (New York: Association for Computing Machinery), doi.org/10.1145/2872518.2890093 (accessed on 1 June 2023).

ONG, Walter.
1982. *Orality and Literacy: The Technologizing of the Word* (New York: Methuen and Co.).

RADFORD, Alec, et al.
2019 “Language models are unsupervised multitask learners”, on *OpenAI* Blog, d4mucfpksywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.PDF (accessed on 1 June 2023).

REBELO, Miguel.
2023 “The best AI chatbots in 2023”, *Zapier.com*, 2 June 2023, zapier.com/blog/best-ai-chatbot/ (accessed on 7 June 2023).

ROBERTSON, Adi.
2023 “I tried the AI novel-writing tool everyone hates, and it’s better than I expected”, *The Verge*, 24 May 2023, www.theverge.com/2023/5/24/23732252/sudowrite-story-engine-ai-generated-cyberpunk-novella (accessed on 4 June 2023).

ROOSE, Kevin.
2022. “An AI-Generated Picture Won an Art Prize. Artists Aren’t Happy”, *The New York Times*, 2 Sept. 2022, www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html (accessed on 1 June 2023).
2022a. “The Brilliance and Weirdness of ChatGPT”, *The New York Times*, 5 Dec. 2022, www.nytimes.com/2022/12/05/technology/chatgpt-ai-twitter.html (accessed on 1 June 2023).
2021. “A Robot wrote this Book Review”, *The New York Times*, 21 Nov. 2021, www.nytimes.com/2021/11/21/books/review/the-age-of-ai-henry-kissinger-eric-schmidt-daniel-huttenlocher.html (accessed on 1 June 2023).
2021. *Futureproof: 9 Rules for Humans in the Age of Automation* (New York: Random House)
2014. “Robots Are Invading the News Business, and It’s Great for Journalists”, *New York Magazine*, 11 July 2014, nymag.com/intelligencer/2014/07/why-robot-journalism-is-great-for-journalists.html (accessed on 1 June 2023).

RUHL, Charlotte.
2020. “Intelligence: Definition, Theories, and Testing”, *Simply Psychology*, 16 July 2020, www.simplypsychology.org/intelligence.html (accessed on 1 June 2023).

SCHATTEN, Jeff.
2022. “Will AI Kill College Writing?”, *The Chronicle of Higher Education*, 14 Sept. 2022, www.chronicle.com/article/will-artificial-intelligence-kill-college-writing (accessed on 1 June 2023).

SCHWEIKARDT, Michael.
2022. “Artificial Intelligence Comes to Theatrical Design,” *The Theatre Times*, 17 Sept. 2022, https://thetheatretimes.com/imagine-artificial-intelligence-comes-to-theatrical-design/ (accessed on 1 June 2023).

TUFEKCI, ZEYNEP.
2022. “What Would Plato Say About ChatGPT?”, *The New York Times*, 15 Dec. 2022, www.nytimes.com/2022/12/15/opinion/chatgpt-education-ai-technology.html (accessed on 6 June 2023).

VYGOTSKY, Lev.
1934. *Thought and Language*, trans. Norris Minick as *The Collected Works of L.S. Vygotsky* (New York: Plenum Press, 1987).