

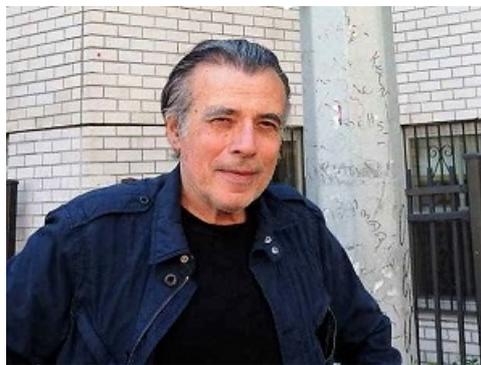


Adventures in writing about creative machines

Arthur I Miller

Opinion - Books Friday, 4th December 2020

Arthur I Miller reports on progress towards the Age of Artificial Superintelligence



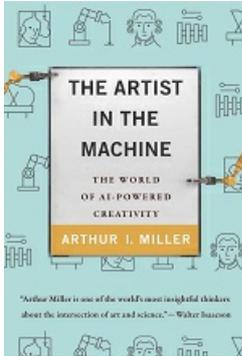
When I began researching my book *The Artist in the Machine* I knew a fair amount about AI-created art and music, but not so much about AI-created literature. From what I did know, I wondered whether I would have enough material to fill a chapter. It turned out that I had more than enough. It turned out, in fact, to be a fascinating and exciting topic, which many AI researchers consider the Final Frontier of AI, in that it involves so many dimensions of intelligence. I was eager to find out what machines themselves might be writing in these very early days - the infancy, one could call it, of computer literacy.

Let's begin with one of the shortest and most deceptively simple forms of writing: jokes. The idea that a computer might tell a joke at all is intriguing in itself. Take the seemingly straightforward: Veni, Vidi, Visa: I came, I saw, I did a little shopping. To dream this up requires, for a start, a knowledge of rudimentary Latin, of Caesar's immortal words, and of what a Visa card is. Could a machine crack such a joke? Would it even know how and when to do so?

How could a machine be made aware of all the necessary nuances and social graces? Would we have to programme them into, say, our laptop's memory? It turns out that laptops belong to a class of computers that are called "ruled-based" machines, or "symbolic machines". They manipulate symbols like words and objects using rules which are also programmed in, and are jam packed with huge databases and sets of rules for dealing with this data. Deep Blue, the chess-playing machine that defeated Garry Kasparov, is an example of a symbolic machine.

Machines that learn

At the other end of the spectrum are artificial neural networks which are inspired by the way the human brain is wired. Their forte is learning. AlphaGo, the algorithm that cracked the game of Go, runs on an artificial neural network. Far from being programmed to play Go, amazingly it taught itself, by studying 30 million board positions from games played by Go masters.



At the moment, neither rule-based machines nor artificial neural networks are up to the task of producing truly funny jokes. Cracking the problem of humour is devilishly difficult, because it's akin to solving the conundrum at the root of AI itself: how to evolve computers as intelligent as human beings?

If we feed an artificial neural network with jokes, such as the most primitive, knock-knock jokes, it will encode the text in numbers, thus learning how such jokes are structured. Using an algorithm that predicts the next word in a sequence, based on what's in its memory, it will create new knock-knock jokes ad infinitum. To start the process, all we have to do is insert the first line of the joke: "Knock-knock."

Here is one such joke created by an artificial neural network: "Knock Knock/Who's there?/Alec/Alec who?/Alec knock knock jokes." So bad it's almost funny, especially since a machine produced it. Conversely, here is a joke created by a symbolic machine programmed to do so: A robot walks into a bar/"What can I get you?" asks the bartender/"I need something to loosen me up," says the robot/So the bartender serves him a screwdriver."

Machine fiction

But could a machine write a story? Symbolic machines, it turns out, can be stocked with thousands of pre-fabricated plots and templates for manipulating them. Their output is prose of a sort that looks familiar, although they have produced nothing really memorable thus far.

But for artificial neural networks it's a whole other ballgame. As they are not rule-based, they have more freedom to write text, to use their imaginations. They can graze on millions of books and articles and, without being programmed to do so, work out how the data is structured. Having programmed it with an algorithm that predicts sequences of words based on what's in its memory, we can give a machine a seeding sentence, from which it will generate text.

I was excited to discover that the computer scientist Ross Goodwin, along with the film maker Oscar Sharpe, had created an AI that wrote a Sci-Fi script. First they fed hundreds of Sci-Fi scripts into their artificial neural network, then gave the machine a seed sentence. From this it generated a Sci-Fi script complete with stage directions which they entitled *Sunspring*. The script is gnomic, to say the least, stream of consciousness taken to the extreme. But it makes sense when spoken by actors with intensity and passion. In one scene the stage direction calls for an actor to spit out his eyeball - no stranger than Shakespeare's famous "Exit pursued by a bear" in *The Winter's Tale*. As Goodwin points out, people who have problems understanding Shakespeare on the printed page do better in the theatre.

Their computer, which they called *Jetson*, was later interviewed and concluded a rambling response to the question, "What's next for you?" with the unexpected and rather touching declaration, "My name is Benjamin." From then on it was Benjamin. Was this assertion just chance or was it a sign of something deeper, a hint of a personality? Was the machine claiming human status?

Sense and nonsense

I am particularly intrigued by work that explores the boundary between sense and what appears at first sight to be nonsense, as in *Sunspring*. AI-created prose is well placed to explore this area in that computers can create text that doesn't work in the way that people expect language to work. In this way AI can change the landscape of language and so the course of literature, just as it has changed the course of art and music.

Will AIs ever create writing that is indistinguishable from the writing we do ourselves? At present the most advanced AI in this field is "Generative Pre-trained Transformer 3", GPT-3 for short, the 3 indicating that it is the third in a series of devices. GPT-3 is an artificial neural network trained on virtually the entire contents of the web, together with blogs and social media. The machine picks up every possible likely connection



between the most far-flung words in its corpus, and can read and write better than any previous system.

AI and fake news

But GPT-3 still has problems. As we might expect it has trouble understanding the meanings of words and is prone to simple errors. More serious, it can be misused. What it writes can be so close to what a person would write that it's sometimes difficult to tell the two apart. This means it can be used to generate misinformation - fake news. Another problem is that since the algorithm was trained on material written by fallible human beings, it is likely to mirror the gender and sexual biases found in society.

Nevertheless GPT-3 is a huge step forward in Natural Language Processing, where the goal is to develop machines that can understand language fluently, with all its nuances and tropes.

The AGI age

When this is achieved machines will be able truly to read and understand the web and in a flash acquire more knowledge than we can in a lifetime. This will pave the way for machines to convince both themselves and us that they are equipped with emotions, consciousness and creativity equivalent to ours. This will be the Age of Artificial General Intelligence.

There is another step after that - the Age of Artificial Superintelligence, when machines will be unimaginably more intelligent than us. By that time what it is to be human too will have dramatically transformed. There will no longer be Artificial Intelligence and Human Intelligence, only Intelligence. To which I will add one thought: are we not already merging with machines?

Arthur I Miller is the author of many critically acclaimed books, including the Pulitzer Prize-nominated *Einstein, Picasso: Space, Time, and the Beauty That Causes Havoc* (Basic Books, 2001); *Insights of Genius: Imagery and Creativity in Science and Art* (MIT Press, 2000); and *Colliding Worlds: How Cutting-Edge Science Is Redefining Contemporary Art* (WW Norton, 2014). He regularly broadcasts, lectures, and curates exhibitions at the intersection of art and science, and has written for the *Guardian*, *New York Times* and *Wired*. He is currently emeritus professor of the history and philosophy of science at University College London. His book on AI and creativity in art, literature and music, *The Artist in the Machine: The World of AI-Powered Creativity*, was published in autumn 2019 by MIT Press, and is just out in paperback.

www.arthurimiller.com; www.collidingworlds.org; www.artistinthemachine.net

Source article: <https://www.bookbrunch.co.uk/page/article-detail/adventures-in-writing-about-creative-machines>