

Can robots write? Machine learning produces dazzling results, but some assembly is still required

September 18 2020, by Alexandra Louise Uitdenbogerd



Credit: AI-generated image (disclaimer)

You might have seen <u>a recent article</u> from The Guardian written by "a robot." Here's a sample:

"I know that my brain is not a 'feeling brain.' But it is capable of making



rational, logical decisions. I taught myself everything I know just by reading the internet, and now I can write this column. My brain is boiling with ideas!"

Read the whole thing and you may be astonished at how coherent and stylistically consistent it is. The software used to produce it is called a "generative model," and they have come a long way in the past year or two.

But exactly *how* was the article created? And is it really true that software "wrote this entire article"?

How machines learn to write

The text was generated using the latest neural network <u>model</u> for language, called <u>GPT-3</u>, released by the American artificial intelligence research company <u>OpenAI</u>. (GPT stands for Generative Pre-trained Transformer.)

OpenAI's previous model, <u>GPT-2</u>, made waves last year. It produced a fairly plausible article about the <u>discovery of a herd of unicorns</u>, and the researchers initially <u>withheld the release</u> of the underlying code for fear it would be abused.

But let's step back and look at what text generation software actually does.

Machine learning approaches fall into three main categories: heuristic models, statistical models, and models inspired by biology (such as <u>neural networks</u> and evolutionary algorithms).

Heuristic approaches are based on "rules of thumb." For example, we learn rules about how to conjugate verbs: I run, you run, he runs, and so



on. These approaches aren't used much nowadays because they are inflexible.

Writing by numbers

Statistical approaches were the state of the art for language-related tasks for many years. At the most basic level, they involve counting words and guessing what comes next.

As a simple exercise, you could generate text by randomly selecting words based on how often they normally occur. About 7% of your words would be "the"—it's the most common word in English. But if you did it without considering context, you might get nonsense like "the the is night aware."

More sophisticated approaches use "bigrams," which are pairs of consecutive words, and "trigrams," which are three-word sequences. This allows a bit of context and lets the current piece of text inform the next. For example, if you have the words "out of," the next guessed word might be "time."

This happens with the auto-complete and auto-suggest features when we write text messages or emails. Based on what we have just typed, what we tend to type and a pre-trained background model, the system predicts what's next.

While bigram- and trigram-based statistical models can produce good results in simple situations, the best recent models go to another level of sophistication: deep learning neural networks.

Imitating the brain



Neural networks work a bit like tiny brains made of several layers of virtual neurons.

A neuron receives some input and may or may not "fire" (produce an output) based on that input. The output feeds into neurons in the next layer, cascading through the network.

The first artificial neuron was <u>proposed in 1943</u> by US neuroscientists Warren McCulloch and Walter Pitts, but they have only become useful for complex problems like generating text in the past five years.

To use neural networks for text, you put words into a kind of numbered index. You can use the number to represent a word, so for example 23,342 might represent "time."

Neural networks do a series of calculations to go from sequences of numbers at the input layer, through the interconnected "hidden layers" inside, to the output layer. The output might be numbers representing the odds for each word in the index to be the next word of the text.

In our "out of" example, number 23,432 representing "time" would probably have much better odds than the number representing "do."

What's so special about GPT-3?

GPT-3 is the latest and best of the text modeling systems, and it's huge. The authors say it has 175 billion parameters, which makes it at least ten times larger than the previous biggest model. The neural network has 96 layers and, instead of mere trigrams, it keeps track of sequences of 2,048 words.

The most expensive and time-consuming part of making a model like this is training it—updating the weights on the connections between



neurons and layers. Training GPT-3 would have used about 262 megawatt-hours of energy, or enough to run my house for 35 years.

GPT-3 can be applied to multiple tasks such as machine translation, autocompletion, answering general questions, and writing articles. While people can often tell its articles are not written by human authors, we are now likely to get it right only <u>about half the time</u>.

The robot writer

But back to how the article in The Guardian was created. GPT-3 needs a prompt of some kind to start it off. The Guardian's staff <u>gave the model</u> <u>instructions</u> and some opening sentences.

This was done eight times, generating eight different articles. The Guardian's editors then combined pieces from the eight generated articles, and "cut lines and paragraphs, and rearranged the order of them in some places," saying "editing GPT-3's op-ed was no different to editing a human op-ed."

This sounds about right to me, based on my own experience with textgenerating software. Earlier this year, my colleagues and I used GPT-2 to write the lyrics for a song we entered in the <u>AI Song Contest</u>, a kind of artificial intelligence Eurovision.

We fine-tuned the GPT-2 model using lyrics from Eurovision songs, provided it with seed words and phrases, then selected the final lyrics from the generated output.

For example, we gave Euro-GPT-2 the seed word "flying," and then chose the output "flying from this world that has gone apart," but not "flying like a trumpet." By automatically matching the lyrics to generated melodies, generating synth sounds based on koala noises, and



applying some great, very human, production work, we got a good result: our song, Beautiful the World, was voted the <u>winner of the contest</u>.

Co-creativity: humans and AI together

So can we really say an AI is an author? Is it the AI, the developers, the users or a combination?

A useful idea for thinking about this is "co-creativity." This means using generative tools to spark new ideas, or to generate some components for our creative work.

Where an AI creates complete works, such as a complete <u>article</u>, the human becomes the curator or editor. We roll our very sophisticated dice until we get a result we're happy with.

This article is republished from <u>The Conversation</u> under a Creative Commons license. Read the <u>original article</u>.

Provided by The Conversation

Citation: Can robots write? Machine learning produces dazzling results, but some assembly is still required (2020, September 18) retrieved 28 April 2024 from <u>https://techxplore.com/news/2020-09-robots-machine-dazzling-results-required.html</u>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.