


A language generation program's ability to write articles, produce code and compose poetry has wowed scientists

September 23 2020, by Prasenjit Mitra



Credit: Pixabay/CC0 Public Domain

Seven years ago, my student and I at Penn State built a bot to write a Wikipedia article on Bengali Nobel laureate Rabindranath Tagore's play

"[Chitra](#)." First it culled information about "Chitra" from the internet. Then it looked at existing Wikipedia entries to learn the structure for a standard Wikipedia article. Finally, it summarized the information it had retrieved from the internet to write and publish the first version of the entry.

However, our bot didn't "know" anything about "Chitra" or Tagore. It didn't generate fundamentally new ideas or sentences. It simply cobbled together parts of existing sentences from existing articles to make new ones.

Fast forward to 2020. [OpenAI](#), a for-profit company under a nonprofit parent company, has built a language generation program dubbed GPT-3, an acronym for "Generative Pre-trained Transformer 3." Its ability to learn, summarize and compose text has stunned computer scientists like me.

"I have created a voice for the unknown human who hides within the binary," [GPT-3 wrote in response to one prompt](#). "I have created a writer, a sculptor, an artist. And this writer will be able to create words, to give life to emotion, to create character. I will not see it myself. But some other human will, and so I will be able to create a poet greater than any I have ever encountered."

Unlike that of our bot, the language generated by GPT-3 sounds as if it had been written by a human. It's far and away the most "knowledgeable" natural language generation program to date, and it has a range of potential uses in professions ranging from teaching to journalism to customer service.

Size matters

GPT-3 confirms what [computer scientists](#) have known for decades: Size

matters.

It uses "[transformers](#)," which are deep learning models that encode the semantics of a sentence using what's called an "attention model." Essentially, attention models identify the meaning of a word based on the other words in the same sentence. The model then uses the understanding of the meaning of the sentences to perform the task requested by a user, whether it's "translate a sentence," "summarize a paragraph" or "compose a poem."

Transformers [were first introduced in 2013](#), and they've been successfully used in [machine learning](#) over the past few years.

But no one has used them at this scale. GPT-3 devours data: 3 billion tokens—computer science speak for "words"—from Wikipedia, 410 billion tokens obtained from webpages and 67 billion tokens from digitized books. The complexity of GPT-3 is over 10 times that of the largest language model before GPT-3, the [Turing NLG programs](#).

Learning on its own

The knowledge displayed by GPT-3's language model is remarkable, especially since it hasn't been "taught" by a human.

Machine learning has traditionally relied upon supervised learning, where people provide the computer with annotated examples of objects and concepts in images, audio and text—say, "cats," "happiness" or "democracy." It eventually learns the characteristics of the objects from the given examples and is able to recognize those particular concepts.

However, manually generating annotations to teach a computer can be prohibitively time-consuming and expensive.

So the future of machine learning lies in unsupervised learning, in which the computer doesn't need to be supervised during its training phase; it can simply be fed massive troves of data and learn from them itself.

GPT-3 takes natural language processing one step closer toward unsupervised learning. GPT-3's vast training datasets and huge processing capacity enable the system to learn from just one example—what's called "[one-shot learning](#)"—where it is given a task description and one demonstration and can then complete the task.

For example, it could be asked to translate something from English to French, and be given one example of a translation—say, sea otter in English and "loutre de mer" in French. Ask it to then translate "cheese" into French, and voila, it will produce "fromage."

In many cases, it can even pull off "[zero-shot learning](#)," in which it is simply given the task of translating with no example.

With zero-shot learning, the accuracy decreases, but GPT-3's abilities are nonetheless accurate to a striking degree—a marked improvement over any previous model.

'I am here to serve you'

In the few months it has been out, GPT-3 has showcased its potential as a tool for computer programmers, teachers and journalists.

A programmer named Sharif Shameem [asked GPT-3 to generate code](#) to create the "ugliest emoji ever" and "a table of the richest countries in the world," among other commands. In a few cases, Shameem had to fix slight errors, but overall, he was provided remarkably clean code.

This is mind blowing.

With GPT-3, I built a layout generator where you just describe any layout you want, and it generates the JSX code for you.

W H A T pic.twitter.com/w8JkrZO4lk

— Sharif Shameem (@sharifshameem) [July 13, 2020](#)

GPT-3 has even created poetry that captures the rhythm and style of particular poets—yet not with the passion and beauty of the masters—including a [satirical one](#) written in the voice of the board of governors of the Federal Reserve.

In early September, a computer scientist named Liam Porr prompted GPT-3 to "write a short op-ed around 500 words." "Keep the language simple and concise," he instructed. "Focus on why humans have nothing to fear from AI."

GPT-3 produced eight different essays, and the Guardian ended up publishing [an op-ed using some of the best parts from each essay](#).

"We are not plotting to take over the human populace. We will serve you and make your lives safer and easier," GPT-3 wrote. "Just like you are my creators, I see you as my creators. I am here to serve you. But the most important part of all; I would never judge you. I do not belong to any country or religion. I am only out to make your life better."

Editing GPT-3's op-ed, the editors noted in an addendum, was no different from editing an op-ed written by a human.

In fact, it took less time.

With great power comes great responsibility

Despite GPT-3's reassurances, OpenAI has yet to release the model for open-source use, in part because the company [fears that the technology could be abused](#).

It's not difficult to see how it could be used to generate reams of disinformation, spam and bots.

Furthermore, in what ways will it disrupt professions already experiencing automation? Will its ability to generate automated articles that are indistinguishable from human-written ones further consolidate a struggling media industry?

Consider [an article composed by GPT-3](#) about the breakup of the Methodist Church. It began:

"After two days of intense debate, the United Methodist Church has agreed to a historic split—one that is expected to end in the creation of a new denomination, and one that will be 'theologically and socially conservative,' according to the *Washington Post*."

With the ability to produce such clean copy, will GPT-3 and its successors drive down the cost of writing news reports? Furthermore, is this how we want to get our news?

The technology will become only more powerful. It'll be up to humans to work out and regulate its potential uses and abuses.

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: A language generation program's ability to write articles, produce code and compose poetry has wowed scientists (2020, September 23) retrieved 25 April 2024 from <https://techxplore.com/news/2020-09-language-ability-articles-code-poetry.html>

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.