


Large computer language models carry environmental, social risks

March 10 2021, by Jackson Holtz



Credit: CC0 Public Domain

Computer engineers at the world's largest companies and universities are using machines to scan through tomes of written material. The goal?

Teach these machines the gift of language. Do that, some even claim, and computers will be able to mimic the human brain.

But this impressive compute capability comes with real costs, including perpetuating racism and causing significant environmental damage, according to a new paper, "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" The paper is being presented Wednesday, March 10 at the ACM Conference on Fairness, Accountability and Transparency (ACM FAccT).

This is the first exhaustive review of the literature surrounding the risks that come with rapid growth of language-learning technologies, said Emily M. Bender, a University of Washington professor of linguistics and a lead author of the paper along with Timnit Gebru, a well-known AI researcher.

"The question we're asking is what are the possible dangers of this approach and the answers that we're giving involve surveying literature across a broad range of fields and pulling them together," said Bender, who is the UW Howard and Frances Nostrand Endowed Professor.

What the researchers surfaced was that there are downsides to the ever-growing computing power put into natural language models. They discuss how the ever-increasing size of training data for language modeling exacerbates social and environmental issues. Alarming, such language models perpetuate hegemonic language and can deceive people into thinking they are having a "real" conversation with a person rather than a machine. The increased computational needs of these models further contributes to environmental degradation.

The authors were motivated to write the paper because of a trend within the field towards ever-larger language models and their growing spheres of influence.

The paper already has generated wide-spread attention due, in part, to the fact that two of the paper's co-authors say they were fired recently from Google for reasons that remain unsettled. Margaret Mitchell and Gebru, the two now-former Google researchers, said they stand by the paper's scholarship and point to its conclusions as a clarion call to industry to take heed.

"It's very clear that putting in the concerns has to happen right now, because it's already becoming too late," said Mitchell, a researcher in AI.

It takes an enormous amount of computing power to fuel the model language programs, Bender said. That takes up energy at tremendous scale, and that, the authors argue, causes environmental degradation. And those costs aren't borne by the [computer engineers](#), but rather by marginalized people who cannot afford the [environmental costs](#).

"It's not just that there's big energy impacts here, but also that the carbon impacts of that will bring costs first to people who are not benefiting from this technology," Bender said. "When we do the [cost-benefit analysis](#), it's important to think of who's getting the benefit and who's paying the cost because they're not the same people."

The large scale of this compute power also can restrict access to only the most well-resourced companies and research groups, leaving out smaller developers outside of the U.S., Canada, Europe and China. That's because it takes huge machines to run the software necessary to make computers mimic human thought and speech.

Another risk comes from the training data itself, the authors say. Because the computers read language from the Web and from other sources, they can pick up and perpetuate racist, sexist, ableist, extremist and other harmful ideologies.

"One of the fallacies that people fall into is well, the internet is big, the internet is everything. If I just scrape the whole internet then clearly I've incorporated diverse viewpoints," Bender said. "But when we did a step-by-step review of the literature, it says that's not the case right now because not everybody's on the internet, and of the people who are on the internet, not everybody is socially comfortable participating in the same way."

And, people can confuse the language models for real human interaction, believing that they're actually talking with a person or reading something that a person has spoken or written, when, in fact, the [language](#) comes from a machine. Thus, the stochastic parrots.

"It produces this seemingly coherent text, but it has no communicative intent. It has no idea what it's saying. There's no there there," Bender said.

More information: Emily M. Bender et al, On the Dangers of Stochastic Parrots, *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (2021). [DOI: 10.1145/3442188.3445922](https://doi.org/10.1145/3442188.3445922)

Provided by University of Washington

Citation: Large computer language models carry environmental, social risks (2021, March 10) retrieved 25 April 2024 from <https://techxplore.com/news/2021-03-large-language-environmental-social.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.
