

Physical training is the next hurdle for artificial intelligence, researcher says

September 27 2022



Credit: Pixabay/CC0 Public Domain



Let a million monkeys clack on a million typewriters for a million years and, the adage goes, they'll reproduce the works of Shakespeare. Give infinite monkeys infinite time, and they still will not appreciate the bard's poetic turn-of-phrase, even if they can type out the words. The same holds true for artificial intelligence (AI), according to Michael Woolridge, professor of computer science at the University of Oxford. The issue, he said, is not the processing power, but rather a lack of experience.

His perspective was published on July 25 in Intelligent Computing.

"Over the past 15 years, the speed of progress in AI in general, and <u>machine learning</u> (ML) in particular, has repeatedly taken seasoned AI commentators like myself by surprise: we have had to continually recalibrate our expectations as to what is going to be possible and when," Wooldridge said.

"For all that their achievements are to be lauded, I think there is one crucial respect in which most large ML models are greatly restricted: the world and the fact that the models simply have no experience of it."

Most ML models are built in virtual worlds, such as video games. They can train on massive datasets, but for physical applications, they are missing vital information. Wooldridge pointed to the AI underpinning autonomous vehicles as an example.

"Letting <u>driverless cars</u> loose on the roads to learn for themselves is a nonstarter, so for this and other reasons, researchers choose to build their models in virtual worlds," Wooldridge said. "And in this way, we are getting excited about a generation of AI systems that simply have no ability to operate in the single most important environment of all: our world."



Language AI models, on the other hand, are developed without a pretense of a world at all—but still suffer from the same limitations. They have evolved, so to speak, from laughably terrible predictive texts to Google's LaMDA, which made headlines earlier this year when a now-former Google engineer claimed the AI was sentient.

"Whatever the validity of [the engineer's] conclusions, it was clear that he was deeply impressed by LaMDA's ability to converse—and with good reason," Wooldridge said, noting that he does not personally believe LaMDA is sentient, nor is AI near such a milestone.

"These foundational models demonstrate unprecedented capabilities in natural language generation, producing extended pieces of naturalsounding text. They also seem to have acquired some competence in common-sense reasoning, one of the holy grails of AI research over the past 60 years."

Such models are <u>neural networks</u>, feeding on enormous datasets and training to understand them. For example, GPT-3, a predecessor to LaMDA, trained on all of the English text available on the internet. The amount of training data combined with significant computing power makes the models akin to <u>human brains</u>, where they move past narrow tasks to begin recognizing patterns and make connections seemingly unrelated to the primary task.

"The bet with foundation models is that their extensive and broad training leads to useful competencies across a range of areas, which can then be specialized for specific applications," Wooldridge said. "While symbolic AI was predicated on the assumption that intelligence is primarily a problem of knowledge, foundation models are predicated on the assumption that intelligence is primarily a problem of data. To simplify, but not by much, throw enough training data at big models, and hopefully, competence will arise."



This "might is right" approach scales the models larger to produce smarter AI, Wooldridge argued, but this ignores the physical know-how needed to truly advance AI.

"To be fair, there are some signs that this is changing," Wooldridge said, pointing to the Gato system. Announced in May by DeepMind, the foundation model, trained on large language sets and on robotic data, could operate in a simple but physical environment.

"It is wonderful to see the first baby steps taken into the physical world by foundation models. But they are just baby steps: the challenges to overcome in making AI work in our world are at least as large—and probably larger—than those faced by making AI work in simulated environments."

More information: Michael Wooldridge, What Is Missing from Contemporary AI? The World, *Intelligent Computing* (2022). <u>DOI:</u> <u>10.34133/2022/9847630</u>

Provided by Intelligent Computing

Citation: Physical training is the next hurdle for artificial intelligence, researcher says (2022, September 27) retrieved 3 May 2024 from <u>https://techxplore.com/news/2022-09-physical-hurdle-artificial-intelligence.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.