


Students develop tool to predict the carbon footprint of algorithms

November 3 2020



Credit: CC0 Public Domain

On a daily basis, and perhaps without realizing it, most of us are in close contact with advanced AI methods known as deep learning. Deep

learning algorithms churn whenever we use Siri or Alexa, when Netflix suggests movies and TV shows based upon our viewing histories, or when we communicate with a website's customer service chatbot.

However, this rapidly evolving technology, one that has otherwise been expected to serve as an effective weapon against [climate change](#), has a downside that many people are unaware of—sky high [energy consumption](#). Artificial intelligence, and particularly the subfield of deep learning, appears likely to become a significant climate culprit should industry trends continue. In only six years—from 2012 to 2018—the computing power needed for deep learning has grown 300,000%. However, the energy [consumption](#) and carbon footprint associated with developing algorithms is rarely measured, despite numerous studies that clearly demonstrate the growing problem.

In response to the problem, two students at the University of Copenhagen's Department of Computer Science, Lasse F. Wolff Anthony and Benjamin Kanding, together with Assistant Professor Raghavendra Selvan, have developed a software program they call Carbontracker. The program can calculate and predict the energy consumption and CO₂ emissions of training deep learning models.

"Developments in this field are going insanely fast and deep learning models are constantly becoming larger in scale and more advanced. Right now, there is exponential growth. And that means an increasing energy consumption that most people seem not to think about," according to Lasse F. Wolff Anthony.

One training session = the annual energy consumption of 126 Danish homes

Deep learning training is the process during which the mathematical

model learns to recognize patterns in large datasets. It's an energy-intensive process that takes place on specialized, power-intensive hardware running 24 hours a day.

"As datasets grow larger by the day, the problems that algorithms need to solve become more and more complex," states Benjamin Kanding.

One of the biggest deep learning models developed thus far is the advanced language model known as GPT-3. In a single [training session](#), it is estimated to use the equivalent of a year's energy consumption of 126 Danish homes, and emit the same amount of CO₂ as 700,000 kilometers of driving.

"Within a few years, there will probably be several models that are many times larger," says Lasse F. Wolff Anthony.

Room for improvement

"Should the trend continue, [artificial intelligence](#) could end up being a significant contributor to climate change. Jamming the brakes on [technological development](#) is not the point. These developments offer fantastic opportunities for helping our climate. Instead, it is about becoming aware of the problem and thinking: How might we improve?" explains Benjamin Kanding.

The idea of Carbontracker, which is a free program, is to provide the field with a foundation for reducing the climate impact of models. Among other things, the program gathers information on how much CO₂ is used to produce energy in whichever region the deep learning training is taking place. Doing so makes it possible to convert energy consumption into CO₂ emission predictions.

Among their recommendations, the two computer science students

suggest that [deep learning](#) practitioners look at when their model trainings take place, as power is not equally green over a 24-hour period, as well as what type of hardware and algorithms they deploy.

"It is possible to reduce the climate impact significantly. For example, it is relevant if one opts to train their model in Estonia or Sweden, where the [carbon footprint](#) of a [model](#) training can be reduced by more than 60 times thanks to greener energy supplies. Algorithms also vary greatly in their energy efficiency. Some require less computing, and thereby less [energy](#), to achieve similar results. If one can tune these types of parameters, things can change considerably," concludes Lasse F. Wolff Anthony.

Provided by University of Copenhagen

Citation: Students develop tool to predict the carbon footprint of algorithms (2020, November 3) retrieved 10 April 2024 from

<https://techxplore.com/news/2020-11-students-tool-carbon-footprint-algorithms.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.
