


Can energy-efficient federated learning save the world?

May 7 2021



Credit: CC0 Public Domain

Training the artificial intelligence models that underpin web search engines, power smart assistants and enable driverless cars consumes megawatts of energy and generates worrying carbon dioxide emissions. But new ways of training these models are proven to be greener.

Artificial intelligence models are used increasingly widely in today's world. Many carry out natural language processing tasks—such as language translation, predictive text and email spam filters. They are also used to empower smart assistants such as Siri and Alexa to "talk" to us, and to operate driverless cars.

But to function well, these models have to be trained on large sets of data, a process that includes carrying out many mathematical operations for every piece of data they are fed. And the data sets they are being trained on are getting ever larger: One recent natural language processing model was trained on a data set of 40 billion words.

As a result, the energy consumed by the [training](#) process is soaring. Most AI models are trained on specialized hardware in large data centers. According to a recent paper in the journal *Science*, the total amount of energy consumed by data centers made up about 1% of [global energy use](#) over the past decade—equalling roughly 18 million US homes. And in 2019, a group of researchers at the University of Massachusetts estimated that training one large AI model used in [natural language](#) processing could generate around the same amount of CO₂ emissions as five cars would generate over their total lifetime.

Concerned by this, researchers at the University of Cambridge set out to investigate more energy-efficient approaches to training AI models. Working with collaborators at the University of Oxford, University College London, and Avignon Université, they explored the environmental impact of a different form of training—called federated learning—and discovered that it had a significantly greener impact. Instead of training the models in data centers, federated learning involves training models across a large number of individual machines. The researchers found that this can lead to lower carbon emissions than traditional learning.

Senior Lecturer Dr. Nic Lane explains how it works when the training is performed not inside large data centers but over thousands of [mobile devices](#)—such as smartphones—where the data is usually collected by the phone users themselves.

"An example of an application currently using federated learning is the next-word prediction in mobile phones," he says. "Each smartphone trains a local model to predict which word the user will type next, based on their previous text messages. Once trained, these local models are then sent to a server. There, they are aggregated into a final model that will then be sent back to all users."

And this method has important privacy benefits as well as environmental benefits, points out Dr. Pedro Porto Buarque De Gusmao, a postdoctoral researcher working with Dr. Lane.

"Users might not want to share the content of their texts with a third party," he explains. "In federated learning, we can keep data local and use the collective power of millions of mobile devices together to train AI models without users' raw data ever leaving the phone."

"And besides these privacy-related gains," says Dr. Lane, "in our recent research, we have shown that federated learning can also have a positive impact in reducing carbon emissions."

"Although smartphones have much less processing power than the hardware accelerators used in data centers, they don't require as much cooling power as the accelerators do. That's the benefit of distributing the training of models across a wide pool of devices."

The researchers recently co-authored a paper on this called "Can Federated Learning save the planet?" and will be discussing their findings at an international research conference, the Flower Summit

2021, on May 11.

In their paper, they offer the first-ever systematic study of the carbon footprint of federated learning. They measured the carbon footprint of a federated learning setup by training two models—one in image classification, the other in speech recognition—using a server and two chipsets popular in the simple devices targeted by federated methods. They recorded the energy consumption during training, and how it might vary depending on where in the world the chipsets and server were located.

They found that while there was a difference between CO₂ emission factors among countries, federated learning under many common application settings was reliably "cleaner" than centralized training.

Training a model to classify images in a large image dataset, they found any federated learning setup in France emitted less CO₂ than any centralized setup in both China and the U.S. And in training the speech recognition model, federated learning was more efficient than centralized training in any country.

Such results are further supported by an expanded set of experiments in a follow-up study ("A first look into the carbon footprint of federated learning") by the same lab that explores an even wider variety of data sets and AI models. And this research also provides the beginnings of necessary formalism and algorithmic foundation of even lower carbon emissions for federated learning in the future.

Based on their research, the researchers have made available a first-of-its-kind "Federated Learning Carbon Calculator" so that the public and other researchers can estimate how much CO₂ is produced by any given pool of devices. It allows users to detail the number and type of devices they are using, which country they are in, which datasets and

upload/download speeds they are using and the number of times each device will train on its own data before sending its [model](#) for aggregation.

They also offer a similar calculator for estimating the carbon emissions of centralized machine learning.

"The development and usage of AI is playing an increasing role in the tragedy that is climate change," says Dr. Lane, "and this problem will only worsen as this technology continues to proliferate through society. We urgently need to address this which is why we are keen to share our findings showing that federated learning methods can produce less CO₂ than data centers under important application scenarios.

"But even more importantly, our research also shines a light as to how federated learning should evolve towards being even more broadly environmentally friendly. Decentralized methods like this will be key in the invention of future sustainable forms of AI in the years ahead."

Provided by University of Cambridge

Citation: Can energy-efficient federated learning save the world? (2021, May 7) retrieved 18 April 2024 from <https://techxplore.com/news/2021-05-energy-efficient-federated-world.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.