

# HOW TO SHRINK AI'S BALLOONING CARBON FOOTPRINT

Emissions data could help researchers to reduce the environmental cost of machine-learning experiments.

By Elizabeth Gibney

**A**s machine-learning experiments get more sophisticated, their carbon footprints are ballooning. Now, researchers have calculated the carbon cost of training a range of models at cloud-computing data centres in various locations. Their findings could help researchers to reduce the emissions created by work that relies on artificial intelligence (AI).

The team found marked differences in emissions between geographical locations. For the same AI experiment, “the most efficient regions produced about a third of the emissions of the least efficient”, says Jesse Dodge, a researcher in machine learning at the Allen Institute for AI in Seattle, Washington, who co-led the study (J. Dodge *et al.* Preprint at <https://arxiv.org/abs/2206.05229>; 2022).

Until now, there haven't been any good tools for measuring emissions produced by cloud-based AI, says Priya Donti, a machine-learning researcher at Carnegie Mellon University in Pittsburgh, Pennsylvania, and co-founder of the group Climate Change AI.

“This is great work done by great authors, and contributes to an important dialogue on how machine-learning workloads can be managed to reduce their emissions,” she says.

Dodge and his collaborators, who included

researchers from Microsoft, monitored electricity consumption while training 11 common AI models, ranging from the kinds of language model that underpin Google Translate to vision algorithms that label images automatically. They put these data together with estimates of how emissions from the electricity grids powering 16 Microsoft Azure cloud-computing servers change over time, to calculate the energy consumption of training in a range of locations.

The team found that training BERT, a common machine-learning language model, at data centres in the central United States or Germany emitted 22–28 kilograms of carbon dioxide, depending on the time of year. This was more than double the emissions generated by doing the same experiment in Norway, which gets most of its electricity from hydroelectric power, or France, which relies mostly on nuclear power (see ‘AI's carbon footprint’).

The time of day at which experiments run also matters. For example, training the AI in Washington during the night, when the state's electricity comes from hydroelectric power alone, led to lower emissions than doing so during the day, when power also comes from gas-fired stations, says Dodge, who presented the results at the Association for Computing Machinery Conference on Fairness, Accountability, and Transparency in Seoul last month.

AI models also varied wildly in their emissions. The image classifier DenseNet created the same CO<sub>2</sub> emissions as charging a mobile phone, whereas training a medium-sized version of a language model known as a transformer (which is much smaller than the popular language model GPT-3, made by research firm OpenAI in San Francisco, California) produced around the same emissions as are generated by a typical US household in a year. Moreover, the team carried out only 13% of the transformer's training process; training it fully would produce emissions “on the order of magnitude of burning an entire railcar full of coal”, says Dodge.

The emissions figures are also underestimates, he adds, because they don't include factors such as the power used for overheads at the data centre, or the emissions that go into creating the necessary hardware. Ideally, the figures would also have included error bars to account for significant underlying uncertainties in a grid's emissions at a given time, says Donti.

## Greener choices

Where other factors are equal, Dodge hopes that the study can help scientists to choose which data centre to use for experiments to minimize emissions. “That decision, it turns out, is one of the most impactful things that someone can do” in the discipline, he says. As a result of the work, Microsoft is now making information on the electricity consumption of its hardware available to researchers who use its Azure service.

Chris Preist at the University of Bristol, UK, who studies the environmental-sustainability impacts of digital technology, says that responsibility for minimizing emissions should lie with the cloud provider rather than the researcher. Providers could ensure that at any one time, the data centres with the lowest carbon intensity are used most, he says. They could also adopt flexible strategies that allow machine-learning runs to start and stop at times that reduce emissions, adds Donti.

Dodge says that the tech companies running the largest experiments should bear the most responsibility for transparency around emissions and for trying to minimize or offset them. Machine learning isn't always bad for the environment, he points out. It can help to design efficient materials, model the climate and track deforestation and endangered species. Nevertheless, the growing carbon footprint of AI is becoming a major cause for concern among some scientists. Even though some research groups are working on tracking carbon outputs, transparency “has yet to grow into something that is the community norm”, says Dodge.

“This work focused on just trying to get transparency on this topic, because that's sorely missing right now,” he says.

## AI'S CARBON FOOTPRINT

The emissions associated with training the language-learning model BERT depend on the time of year, and on the location of the data centre.

