


Commercial cloud service providers give artificial intelligence computing a boost

September 17 2019, by Kim Martineau



MIT researchers are training a pair of generative adversarial networks, or GANs, to mimic the land, sea, and cloud textures seen in satellite images with the goal of eventually visualizing real-world sea-level rise. It's one of many artificial intelligence research projects made possible by IBM and Google-donated cloud credits. Credit: Brandon Leshchinskiy

Neural networks have given researchers a powerful tool for looking into the future and making predictions. But one drawback is their insatiable need for data and computing power ("compute") to process all that information. At MIT, demand for compute is estimated to be five times greater than what the Institute can offer. To help ease the crunch, industry has stepped in. An \$11.6 million supercomputer recently donated by IBM comes online this fall, and in the past year, both IBM and Google have provided cloud credits to MIT Quest for Intelligence for distribution across campus. Four projects made possible by IBM and Google cloud donations are highlighted below.

Smaller, faster, smarter neural networks

To recognize a cat in a picture, a [deep learning model](#) may need to see millions of photos before its artificial neurons "learn" to identify a cat. The process is computationally intensive and carries a steep environmental cost, as new research attempting to measure artificial intelligence's (AI's) carbon footprint has highlighted.

But there may be a more efficient way. New MIT research shows that models only a fraction of the size are needed. "When you train a big network there's a small one that could have done everything," says Jonathan Frankle, a [graduate student](#) in MIT's Department of Electrical Engineering and Computer Science (EECS).

With study co-author and EECS Professor Michael Carbin, Frankle estimates that a [neural network](#) could get by with one-tenth the number of connections if the right subnetwork is found at the outset. Normally, neural networks are trimmed after the training process, with irrelevant connections removed then. Why not train the small model to begin with, Frankle wondered?

Experimenting with a two-neuron network on his laptop, Frankle got

encouraging results and moved to larger image-datasets like MNIST and CIFAR-10, borrowing GPUs where he could. Finally, through IBM Cloud, he secured enough compute power to train a real ResNet model. "Everything I'd done previously was toy experiments," he says. "I was finally able to run dozens of different settings to make sure I could make the claims in our paper."

Frankle spoke from Facebook's offices, where he worked for the summer to explore ideas raised by his Lottery Ticket Hypothesis paper, one of two picked for a best paper award at this year's International Conference on Learning Representations. Potential applications for the work go beyond image classification, Frankle says, and include reinforcement learning and natural language processing models. Already, researchers at Facebook AI Research, Princeton University, and Uber have published follow-on studies.

"What I love about neural networks is we haven't even laid the foundation yet," says Frankle, who recently shifted from studying cryptography and tech policy to AI. "We really don't understand how it learns, where it's good and where it fails. This is physics 1,000 years before Newton."

Distinguishing fact from fake news

Networking platforms like Facebook and Twitter have made it easier than ever to find quality news. But too often, real news is drowned out by misleading or outright false information posted online. Confusion over a recent video of U.S. House Speaker Nancy Pelosi doctored to make her sound drunk is just the latest example of the threat misinformation and fake news pose to democracy.

"You can put just about anything up on the internet now, and some people will believe it," says Moin Nadeem, a senior and EECS major at

MIT.

If technology helped create the problem, it can also help fix it. That was Nadeem's reason for picking a superUROP project focused on building an automated system to fight fake and misleading news. Working in the lab of James Glass, a researcher at MIT's Computer Science and Artificial Intelligence Laboratory, and supervised by Mitra Mohtarami, Nadeem helped train a language model to fact-check claims by searching through Wikipedia and three types of news sources rated by journalists as high-quality, mixed-quality or low-quality.

To verify a claim, the model measures how closely the sources agree, with higher agreement scores indicating the claim is likely true. A high disagreement score for a claim like, "ISIS infiltrates the United States," is a strong indicator of fake news. One drawback of this method, he says, is that the model doesn't identify the independent truth so much as describe what most people think is true.

With the help of Google Cloud Platform, Nadeem ran experiments and built an interactive website that lets users instantly assess the accuracy of a claim. He and his co-authors presented their results at the North American Association of Computational Linguistics (NAACL) conference in June and are continuing to expand on the work.

"The saying used to be that seeing is believing," says Nadeem, in this video about his work. "But we're entering a world where that isn't true. If people can't trust their eyes and ears it becomes a question of what can we trust?"

Visualizing a warming climate

From rising seas to increased droughts, the effects of climate change are already being felt. A few decades from now, the world will be a warmer,

drier, and more unpredictable place. Brandon Leshchinskiy, a graduate student in MIT's Department of Aeronautics and Astronautics (AeroAstro), is experimenting with generative adversarial networks, or GANs, to imagine what Earth will look like then.

GANs produce hyper-realistic imagery by pitting one neural network against another. The first network learns the underlying structure of a set of images and tries to reproduce them, while the second decides which images look implausible and tells the first network to try again.

Inspired by researchers who used GANs to visualize sea-level rise projections from street-view images, Leshchinskiy wanted to see if satellite imagery could similarly personalize climate projections. With his advisor, AeroAstro Professor Dava Newman, Leshchinskiy is currently using free IBM Cloud credits to train a pair of GANs on images of the eastern U.S. coastline with their corresponding elevation points. The goal is to visualize how sea-level rise projections for 2050 will redraw the coastline. If the project works, Leshchinskiy hopes to use other NASA datasets to imagine future ocean acidification and changes in phytoplankton abundance.

"We're past the point of mitigation," he says. "Visualizing what the world will look like three decades from now can help us adapt to climate change."

Identifying athletes from a few gestures

A few moves on the field or court are enough for a computer vision [model](#) to identify individual athletes. That's according to preliminary research by a team led by Katherine Gallagher, a researcher at MIT Quest for Intelligence.

The team trained computer vision models on video recordings of tennis

matches and soccer and basketball games and found that the models could recognize individual players in just a few frames from key points on their body providing a rough outline of their skeleton.

The team used a Google Cloud API to process the video data, and compared their models' performance against models trained on Google Cloud's AI platform. "This pose information is so distinctive that our models can identify players with accuracy almost as good as models provided with much more information, like hair color and clothing," she says.

Their results are relevant for automated player identification in sports analytics systems, and they could provide a basis for further research on inferring player fatigue to anticipate when players should be swapped out. Automated pose detection could also help athletes refine their technique by allowing them to isolate the precise moves associated with a golfer's expert drive or a tennis player's winning swing.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Commercial cloud service providers give artificial intelligence computing a boost (2019, September 17) retrieved 10 April 2024 from <https://techxplore.com/news/2019-09-commercial-cloud-artificial-intelligence-boost.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.
