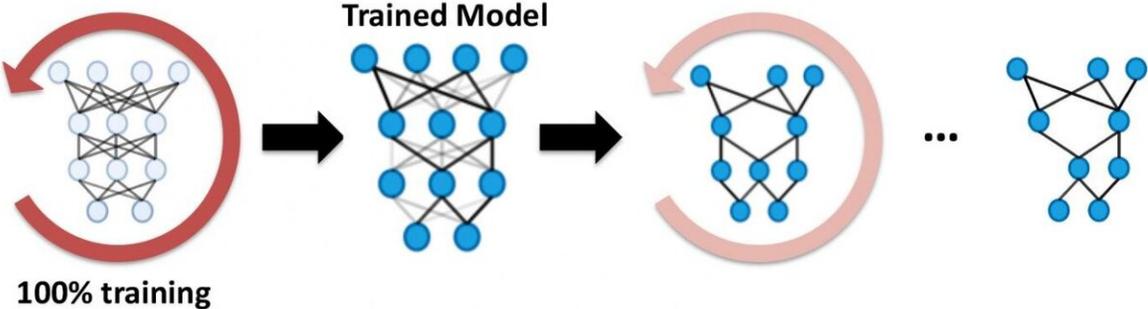


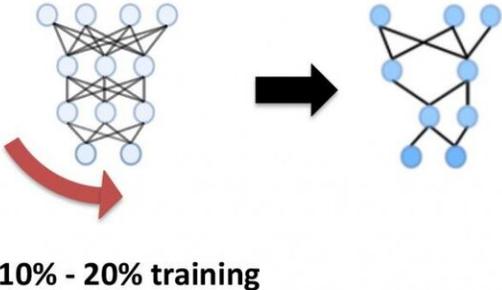
Early Bird uses 10 times less energy to train deep neural networks

May 18 2020, by Jade Boyd

Progressive Pruning and Training



Early-Bird Train (Proposed)



Rice University's Early Bird method for training deep neural networks finds key connectivity patterns early in training, reducing the computations and carbon footprint for the increasingly popular form of artificial intelligence known as deep learning. (Graphic courtesy of Y. Lin/Rice University) Credit: Y. Lin/Rice

University

Rice University's Early Bird could care less about the worm; it's looking for megatons of greenhouse gas emissions.

Early Bird is an energy-efficient method for training deep neural networks (DNNs), the form of artificial intelligence (AI) behind self-driving cars, intelligent assistants, facial recognition and dozens more high-tech applications.

Researchers from Rice and Texas A&M University [unveiled Early Bird](#) April 29 in a spotlight paper at ICLR 2020, the International Conference on Learning Representations. A study by lead authors Haoran You and Chaojian Li of Rice's Efficient and Intelligent Computing (EIC) Lab showed Early Bird could use 10.7 times less energy to train a DNN to the same level of accuracy or better than typical training. EIC Lab director Yingyan Lin led the research along with Rice's Richard Baraniuk and Texas A&M's Zhangyang Wang.

"A major driving force in recent AI breakthroughs is the introduction of bigger, more expensive DNNs," Lin said. "But training these DNNs demands considerable energy. For more innovations to be unveiled, it is imperative to find 'greener' training methods that both address [environmental concerns](#) and reduce financial barriers of AI research."

Training cutting-edge DNNs is costly and getting costlier. A [2019 study](#) by the Allen Institute for AI in Seattle found the number of computations needed to train a top-flight [deep neural network](#) increased 300,000 times between 2012-2018, and [a different 2019 study](#) by researchers at the University of Massachusetts Amherst found the [carbon footprint](#) for training a single, elite DNN was roughly equivalent

to the lifetime carbon dioxide emissions of five U.S. automobiles.

DNNs contain millions or even billions of artificial neurons that learn to perform specialized tasks. Without any explicit programming, deep networks of artificial neurons can learn to make humanlike decisions—and even outperform human experts—by "studying" a large number of previous examples. For instance, if a DNN studies photographs of cats and dogs, it learns to recognize cats and dogs. AlphaGo, a deep network trained to play the board game Go, beat a professional human player in 2015 after studying tens of thousands of previously played games.

"The state-of-art way to perform DNN training is called progressive prune and train," said Lin, an assistant professor of electrical and computer engineering in Rice's Brown School of Engineering. "First, you train a dense, giant network, then remove parts that don't look important—like pruning a tree. Then you retrain the pruned network to restore performance because performance degrades after pruning. And in practice you need to prune and retrain many times to get good performance."

Pruning is possible because only a fraction of the artificial neurons in the network can potentially do the job for a specialized task. Training strengthens connections between necessary neurons and reveals which ones can be pruned away. Pruning reduces model size and computational cost, making it more affordable to deploy fully trained DNNs, especially on small devices with limited memory and processing capability.

"The first step, training the dense, giant network, is the most expensive," Lin said. "Our idea in this work is to identify the final, fully functional pruned network, which we call the 'early-bird ticket,' in the beginning stage of this costly first step."

By looking for key [network](#) connectivity patterns early in training, Lin and colleagues were able to both discover the existence of early-bird tickets and use them to streamline DNN training. In experiments on various benchmarking data sets and DNN models, Lin and colleagues found Early Bird could emerge as little as one-tenth or less of the way through the initial phase of training.

"Our method can automatically identify early-bird tickets within the first 10% or less of the training of the dense, giant networks," Lin said. "This means you can train a DNN to achieve the same or even better accuracy for a given task in about 10% or less of the time needed for traditional [training](#), which can lead to more than one order savings in both computation and energy."

Developing techniques to make AI greener is the main focus of Lin's group. Environmental concerns are the primary motivation, but Lin said there are multiple benefits.

"Our goal is to make AI both more environmentally friendly and more inclusive," she said. "The sheer size of complex AI problems has kept out smaller players. Green AI can open the door enabling researchers with a laptop or limited computational resources to explore AI innovations."

Provided by Rice University

Citation: Early Bird uses 10 times less energy to train deep neural networks (2020, May 18) retrieved 20 September 2024 from <https://techxplore.com/news/2020-05-early-bird-energy-deep-neural.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.