Collaborative machine learning that preserves privacy

7 September 2022, by Adam Zewe

Credit: Pixabay/CC0 Public Domain

Training a machine-learning model to effectively perform a task, such as image classification, involves showing the model thousands, millions, or even billions of example images. Gathering such enormous datasets can be especially challenging when privacy is a concern, such as with medical images. Researchers from MIT and the MIT-born startup DynamoFL have now taken one popular solution to this problem, known as federated learning, and made it faster and more accurate.

Federated learning is a collaborative method for training a machine-learning model that keeps sensitive user data private. Hundreds or thousands of users each train their own model using their own data on their own device. Then users transfer their models to a central server, which combines them to come up with a better model that it sends back to all users.

A collection of hospitals located around the world, for example, could use this method to train a machine-learning model that identifies brain tumors in medical images, while keeping patient data secure on their local servers.

But federated learning has some drawbacks. Transferring a large machine-learning model to and from a central server involves moving a lot of data, which has high communication costs, especially since the model must be sent back and forth dozens or even hundreds of times. Plus, each user gathers their own data, so those data don't necessarily follow the same statistical patterns, which hampers the performance of the combined model. And that combined model is made by taking an average—it is not personalized for each user.

The researchers developed a technique that can simultaneously address these three problems of federated learning. Their method boosts the accuracy of the combined machine-learning model while significantly reducing its size, which speeds up communication between users and the central server. It also ensures that each user receives a model that is more personalized for their environment, which improves performance.

The researchers were able to reduce the model size by nearly an order of magnitude when compared to other techniques, which led to communication costs that were between four and six times lower for individual users. Their technique was also able to increase the model's overall accuracy by about 10 percent.

"A lot of papers have addressed one of the problems of federated learning, but the challenge was to put all of this together. Algorithms that focus just on personalization or communication efficiency don't provide a good enough solution. We wanted to be sure we were able to optimize for everything, so this technique could actually be used in the real world," says Vaikkunth Mugunthan Ph.D. '22, lead author of a paper that introduces this technique.

Mugunthan wrote the paper with his advisor, senior author Lalana Kagal, a principal research scientist in the Computer Science and Artificial Intelligence Laboratory (CSAIL). The work will be presented at

the European Conference on Computer Vision.

Cutting a model down to size

The system the researchers developed, called FedLTN, relies on an idea in machine learning known as the lottery ticket hypothesis. This hypothesis says that within very large neural network models there exist much smaller subnetworks that can achieve the same performance. Finding one of these subnetworks is akin to finding a winning lottery ticket. (LTN stands for "lottery ticket network.")

Neural networks, loosely based on the human brain, are machine-learning models that learn to solve problems using interconnected layers of nodes, or neurons.

Finding a winning lottery ticket network is more complicated than a simple scratch-off. The researchers must use a process called iterative pruning. If the model's accuracy is above a set threshold, they remove nodes and the connections between them (just like pruning branches off a bush) and then test the leaner neural network to see if the accuracy remains above the threshold.

Other methods have used this pruning technique for federated learning to create smaller machine-learning models which could be transferred more efficiently. But while these methods may speed things up, model performance suffers.

Mugunthan and Kagal applied a few novel techniques to accelerate the pruning process while making the new, smaller models more accurate and personalized for each user.

They accelerated pruning by avoiding a step where the remaining parts of the pruned neural network are "rewound" to their original values. They also trained the model before pruning it, which makes it more accurate so it can be pruned at a faster rate, Mugunthan explains.

To make each model more personalized for the user's environment, they were careful not to prune away layers in the network that capture important statistical information about that user's specific data. In addition, when the models were all combined, they made use of information stored in the central server so it wasn't starting from scratch for each round of communication.

They also developed a technique to reduce the number of communication rounds for users with resource-constrained devices, like a smart phone on a slow network. These users start the federated learning process with a leaner model that has already been optimized by a subset of other users.

Winning big with lottery ticket networks

When they put FedLTN to the test in simulations, it led to better performance and reduced communication costs across the board. In one experiment, a traditional federated learning approach produced a model that was 45 megabytes in size, while their technique generated a model with the same accuracy that was only 5 megabytes. In another test, a state-of-the-art technique required 12,000 megabytes of communication between users and the server to train one model, whereas FedLTN only required 4,500 megabytes.

With FedLTN, the worst-performing clients still saw a performance boost of more than 10 percent. And the overall model accuracy beat the state-of-the-art personalization algorithm by nearly 10 percent, Mugunthan adds.

Now that they have developed and finetuned FedLTN, Mugunthan is working to integrate the technique into a federated learning startup he recently founded, DynamoFL.

Moving forward, he hopes to continue enhancing this method. For instance, the researchers have demonstrated success using datasets that had labels, but a greater challenge would be applying the same techniques to unlabeled data, he says.

Mugunthan is hopeful this work inspires other researchers to rethink how they approach federated learning.

"This work shows the importance of thinking about these problems from a holistic aspect, and not just

individual metrics that have to be improved. Sometimes, improving one metric can actually cause a downgrade in the other metrics. Instead, we should be focusing on how we can improve a bunch of things together, which is really important if it is to be deployed in the real world," he says.

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