SCIENCE · PASSION · TECHNOLOGY

Open Thesis / Project

Distributed Training With Local Updates

(Embedded Information Processing Team)

Motivation

Distributed machine learning allows model training on decentralised data residing on various devices, such as mobile phones or IoT devices. However, all these edge devices usually have limited communication bandwidth to transfer the initial global model and local gradient updates. Limited bandwidth is one of the major bottlenecks that hinder applying FL in practice. We would like to leverage permutation invariance of neural networks to make these devices start learning from locally initialised, rather than a global model and enable partial parameter updates. This approach exploits update locality and should considerably reduce bandwidth usage. The goal of the thesis would be to implement the approach and compare it with state-of-the-art distributed machine learning implementations. See https://drive.google.com/file/d/1nOQ8pY1p7h6P9Ai8eweHbv54-VxXTR8R/view for more information. **Interested? Contact us for more details!**

Target Group

Students in ICE or Computer Science.

Thesis Type

Master Project / Master Thesis.

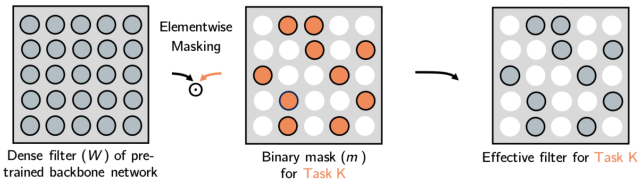

Image source: https://arxiv.org/abs/1801.06519

We will use several tricks, such as supermasks, neuron alignment, permutation invarience to minimize bandwidth usage. You will be provided a starting point to understand these these optimizations.

Goals and Tasks

In this project, you will enable efficient distributed model training without having nodes to share the same initial model and by communicating only local updates. You will be given a solid starting point which supports this optimization. Your task will be as follows:

- Literature review on distributed and federated training, training and aggregating models from different initializations and partial model update;
- Implement distributed model training with provided optimizations;
- Compare the obtained performance to vanilla baselines (the code will be available) and to a non-distributed training;
- Report obtained results in a written report, oral presentation.

Requirements / Skills:

- Good knowledge of deep neural networks and interest in optimization and distributed training;
- Programming skills in Python;
- Prior experience in deep learning frameworks is desirable (preferably PyTorch).

Used Tools / Equipment:

- A laptop (GPU infrastructure will be provided if needed);
- Your talent (very important!).

Contact Person

- Rahim Entezari (entezari@tugraz.at)
- Dr. Olga Saukh (saukh@tugraz.at)

Institute for Technical Informatics
Networked Embedded Systems Group