Abstract Title:

Personalized Federated Learning with Gaussian Processes

Abstract Body:

In personal federated learning (PFL), a server node stores a global model and connects to multiple clients, which have private data that cannot be shared. The goal is to learn the global model in a communication-efficient manner.

We put forward pFedGP, a solution to PFL that is based on Gaussian processes (GPs) with deep kernel learning [1]. We propose learning a shared kernel function across all clients, parameterized by a neural network, with a personal GP classifier for each client. We further extend pFedGP to two settings, the first helps to improve generalization in the low data regime and the second reduces the computational cost. Extensive experiments on standard PFL benchmarks, and on a new setup of learning under input noise show that pFedGP achieves well-calibrated predictions while significantly outperforming baseline methods, reaching up to 21% in accuracy gain.

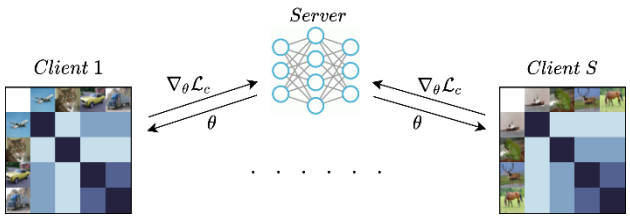

Figure 1: pFedGP. The data at each client is first mapped to an embedding space using a shared neural network across all clients. Then, a GP is applied to the data of the client for model learning and inference.

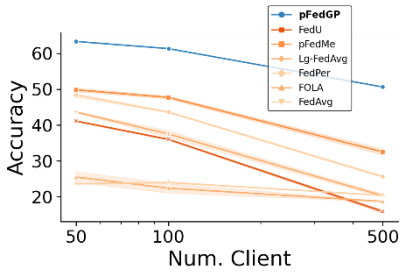

Figure 2: Test accuracy (±SEM) over 50, 100, 500 clients on CIFAR-100.

References:

[1] Wilson, A. G., Hu, Z., Salakhutdinov, R., & Xing, E. P. (2016, May). Deep kernel learning. In Artificial intelligence and statistics (pp. 370-378). PMLR.