

Federated Uncertainty Quantification: a Survey

Speaker: Eric Moulines, École polytechnique



Abstract

Many machine learning applications require training a centralized model on decentralized, heterogeneous, and potentially private data sets. Federated learning (FL) has emerged as a privacy-friendly training paradigm that does not require clients' private data to leave their local devices. FL brings new challenges in addition to "traditional" distributed learning: expensive communication, statistical heterogeneity, partial participation, and privacy.

The "classical" formulation of FL treats it as a distributed optimization problem. Yet the standard distributed optimization algorithms (e.g., data-parallel SGD) are too communication-intensive to be practical at FL. An alternative approach is to consider a Bayesian formulation of the FL problem. Typically within this approach an exact posterior inference is untractable even for models and data sets of modest size, thus an approximate inference methods should be considered. Among the many proposed approaches, we will discuss the MCMC solution, which is the Federated Averaging Langevin Dynamics. We will also cover the approach, based on variational inference, where fewer lockstep synchronization and communication steps may be required between clients and servers.

