

Personalization by Combination

The promises and problems of multi-armed bandit orchestrated federated learning

Fredrika Lundahl*

June 2021

Abstract

This thesis is written in collaboration with Ericsson AI Research and is investigating the statistical aspects of a novel, yet unpublished, method to conduct personalized federated learning. The method is the result of a search for a communication efficient personalization algorithm which is also good at handling poisonous workers.

The proposed model originates from the Federated Averaging algorithm but replaces the averaging over the models of all (or most) workers by the averaging over the models from a selected combination of workers only. The combination is selected by a multi-armed bandit inspired decision rule, which learns the appropriate set of workers to combine in order to optimize the performance on a given target worker. Consequently, the decision rule also learns which workers to avoid.

The model is the most advantageous when the number of workers is small and each worker has a small dataset, situations where other methods struggle. However, having a reinforcement learning model inside a federated learning model causes some challenges, mainly relating to variance caused by the federated learning model. Some ways to combat these challenges are proposed.

By incorporating cluster analysis in future works, we might be able to give the model wider applications as it then should be able to personalize successfully even when the number of workers is big and each worker contains more data.

*Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden.
E-mail: ffredrikaa@gmail.com. Supervisor: Chun-Biu Li.