Federated learning for digital pathology: training algorithms without accessing patient data to protect patient privacy

Fahime Sheikhzadeh[1], Ishan Shah[1], Srividya Doss[1], Yao Nie[1]
[1] Digital Pathology, Roche Diagnostics, Santa Clara, CA, USA

Introduction: Herein, we explore the development of a Federated Learning (FL) system for Digital Pathology (DP) applications. Federated learning as a decentralized training approach enables different hospitals to collaboratively train a machine learning model while keeping their patient data local. In this approach, multiple client devices collaborate and train a deep learning model (global model) without sharing their training data. A server maintains and updates the global model in an iterative process. At each iteration, each client receives and retrains the global model on their local data, and sends it back to the server. The models from all clients get aggregated at the server site. The aggregated model is then sent back to the clients. The iterations continue until training converges.

Methods: In this study, we developed a U-Net based deep learning model for tumor, stroma, and other tissue type segmentation on H&E images. In total, 221 images with the approximate size of 1024 by 1024 pixels, at 10X magnification were extracted from 77 whole slide images and used in our experiments. A fixed dataset at the server was used for validation. Models that were generated at the client site are not controlled by the server and were shared with the server only based on the client privacy policy. Each client had an independent validation dataset and used it to examine the performance of the model based on their quality standards. Based on this validation, the client decided to deploy the global model or not. We performed three experiments including training the model with three clients, with Independent and Identically Distributed (IID) data, three clients with non-IID data, and six clients with partial participation in the FL process.

Results: We successfully trained the segmentation model in a federated manner, without accessing the client data. In our experiments, the global model outperformed the model trained on the centralized data or had a comparable performance.

Conclusion: Establishing a FL system enables us to capitalize on real world data that should stay at hospitals or laboratories, and train a global model that is suitable for all the users to make better diagnostic decisions.