Balancing Privacy and Performance in Federated Learning with Adaptive Differential Privacy and Secure Multi-Party Computation

T. D. Attygalle, A. Athukorala

University of Colombo School of Computing, Sri Lanka

Federated Learning (FL) enables collaborative model training without sharing raw data, but exchanged updates remain vulnerable to inference attacks. While Differential Privacy (DP) adds noise to protect privacy, client-side noise in methods like Adaptive DP-FL introduces high variance and relies on server trust. Secure Multi-Party Computation (SMPC) safeguards updates via secret sharing but lacks integration with differential privacy. We propose a hybrid FL framework combining adaptive DP-FL with SMPC. Clients clip gradients adaptively and secret-share updates across non-colluding servers. Servers reconstruct only the global sum, add a single calibrated Gaussian noise term, and account privacy using R'enyi DP. This design reduces sensitivity by $1/K$, cuts noise variance, and streamlines privacy accounting. Experiments on MNIST and Fashion-MNIST under varying budgets show improved accuracy over the baseline, demonstrating enhanced privacy–utility trade-offs for sensitive applications.

Keywords: *Federated Learning, Differential Privacy, Secure Multi-Party Computation, Shamir's Secret Sharing, R'enyi Differential Privacy*