Synthesis

Eros Fanì

June 2025

Federated Learning (FL) has emerged as a promising approach for training machine learning models using decentralized data sources while maintaining users' privacy. Unlike traditional centralized learning, FL allows users to train models collaboratively without needing to collect their raw data on a central server, making FL particularly suitable for privacy-sensitive applications such as healthcare, finance, and autonomous systems. However, FL faces several challenges, such as statistical heterogeneity, which arises in the presence of variations in data distributions across clients. In computer vision applications, factors such as class and domain imbalance are primary sources of heterogeneity, leading to notable performance degradation and instability during federated training.

This manuscript addresses the challenges of statistical heterogeneity in FL, particularly in relation to computer vision tasks. First, it introduces a novel FL algorithm, Federated Recursive Ridge Regression (`Fed3R`), designed to mitigate the effects of heterogeneity. `Fed3R` constructs a ridge regression-based classifier equivalent to the centralized ridge regression solution. The parameters obtained from `Fed3R` can be used to initialize the classifier for further fine-tuning with other FL algorithms, significantly improving their efficiency and robustness. Then, the importance of initialization in Personalized Federated Learning (PFL) is explored, demonstrating how well-initialized classifiers can substantially improve the performance and efficiency of personalized models. Furthermore, a new method is introduced, Only Local Labels (`OLL`), which enhances local personalization by selectively filtering classifier neurons.

Additionally, FL is studied in the context of autonomous driving for the semantic segmentation task. A new benchmark, `FedDrive`, is introduced to analyze the impact of domain shift and class imbalance in this scenario. Furthermore, a new FL problem, Federated Source-Free Domain Adaptation `FFreeDA`, is proposed, addressing scenarios where clients lack labeled data. Finally, a novel algorithm, Learning Across Domains and Devices (`LADD`), is designed to address (`FFreeDA`), leveraging semi-supervised learning and clustered FL techniques.

Extensive experimental evaluations validate the proposed methods, highlighting their effectiveness in addressing statistical heterogeneity in FL. By ad-

vancing the understanding and practical methodologies for FL for computer vision applications, this manuscript contributes to the broader goal of making FL more robust, scalable, and applicable to real-world scenarios. The findings provide insights into the interplay between FL algorithms, statistical heterogeneity, and computer vision tasks, paving the way for future research in this rapidly evolving field.