FedGuCci: Making Local Models More Connected in Landscape for Federated Learning at Edge

Zexi Li*[1], Dongqi Cai†[1], and Nicholas D. Lane‡[1]

[1]University of Cambridge

Abstract

Federated Learning (FL) presents a powerful paradigm for privacy-preserving, collaborative model training on decentralized edge devices. However, a critical challenge hindering its widespread adoption is the significant generalization gap compared to centralized training. This problem is particularly acute in mobile and edge settings due to extreme data heterogeneity—where each device's local dataset is non-IID—causing locally trained models to drift into disparate regions of the parameter space. Standard aggregation methods like FedAvg consequently fail to fuse these diverged models into a high-performing global model. This paper addresses this fundamental issue from the perspective of "connectivity" in the loss landscape, drawing inspiration from Linear Mode Connectivity (LMC).

LMC studies the path between two neural network solutions, where a high "loss barrier" along their linear interpolation signifies they are in different basins. This phenomenon is analogous to model drift in FL. We introduce and formalize a novel property: the **transitivity of LMC**. We hypothesize that if two models, w_1 and w_2, are independently trained to have a low-loss connection (i.e., a low barrier) to a common, fixed anchor model (w^*_{anc}), their own mutual connectivity is consequently improved. We provide both theoretical proofs and empirical evidence to validate this transitivity, showing it extends from pairs to the "group connectivity" of multiple models, a scenario central to FL.

Building on this principle, we propose **FedGuCci**, a novel client-side FL algorithm designed to improve the group connectivity of local models at the edge. **FedGuCci** cleverly utilizes one or more historical global models as shared anchors. During local training, each client minimizes not only its local task loss but also a *connectivity loss*, which encourages its model to stay in a region of the landscape that is well-connected to these anchors. This is achieved by penalizing high loss values along the linear path to the anchor(s). As this is a purely client-side modification, it requires no extra communication, making it ideal for bandwidth-constrained edge networks.

To further enhance performance under severe heterogeneity, we introduce **FedGuCci+**. This strengthened version explicitly aligns the varied local loss landscapes by incorporating two heterogeneity-resistant modules: (1) Logit Calibration to mitigate the bias caused by label distribution skew, and (2) Sharpness-Aware Minimization (SAM) to guide local optimizers toward flatter minima. Flatter minima naturally create larger, more overlapping low-loss regions, further improving connectivity.

We conducted extensive experiments validating our methods across a diverse array of tasks, from vision (Fashion-MNIST, CIFAR-10/100, Tiny-ImageNet) to natural language understanding (6 GLUE benchmark tasks). **FedGuCci** and **FedGuCci+** consistently achieve state-of-the-art generalization, demonstrating robust performance gains especially under challenging non-IID settings and with a large number of local training epochs. The approach also excels when finetuning large pretrained vision (ViT, ResNet-18) and language (RoBERTa) models, highlighting its potential for collaboratively adapting foundation models on edge devices. By fundamentally improving how local models relate to one another in the parameter landscape, **FedGuCci** offers a communication-efficient and highly effective solution for building generalized and robust AI models at the mobile edge.

*zexi.li@zju.edu.cn
†dc912@cam.ac.uk
‡ndl32@cam.ac.uk