
Federated Learning over Connected Modes

Dennis Grinwald^{1,2}, Philipp Wiesner², Shinichi Nakajima^{1,2,3}

¹BIFOLD, ²TU Berlin, ³RIKEN Center for AIP

{dennis.grinwald, wiesner, nakajima}@tu-berlin.de

Abstract

Statistical heterogeneity in federated learning poses two major challenges: slow global training due to conflicting gradient signals, and the need of personalization for local distributions. In this work, we tackle both challenges by leveraging recent advances in *linear mode connectivity* — identifying a linearly connected low-loss region in the weight space of neural networks, which we call solution simplex. We propose federated learning over connected modes (FLOCO), where clients are assigned local subregions in this simplex based on their gradient signals, and together learn the shared global solution simplex. This allows personalization of the client models to fit their local distributions within the degrees of freedom in the solution simplex and homogenizes the update signals for the global simplex training. Our experiments show that FLOCO accelerates the global training process, and significantly improves the local accuracy with minimal computational overhead.

1 Introduction

Federated learning (FL) [1] is a decentralized machine learning paradigm that facilitates collaborative model training across distributed devices while preserving data privacy. However, in typical real applications, statistical heterogeneity—non-identically and independently distributed (non-IID) data distributions at clients—makes it difficult to train well-performing models. To tackle this difficulty, various methods have been proposed, e.g., personalized FL [2], clustered FL [3], advanced client selection strategies [4], robust aggregation [5], and federated meta- and multi-task learning approaches [6]. These methods aim either at training a global model that performs well on the global distribution [7], or, as it is common in personalized FL, at training multiple client-dependent models each of which performs well on its local distribution [8]. These two aims often pose a trade-off—a model that shows better global performance tends to show worse local performance, and vice versa. In this work, we aim to develop a FL method that improves both global and local performance.

Our approach leverages recent findings on *mode connectivity* [9–11]—the existence of low-loss paths in the weight parameter space between independently trained neural networks—and its applications [12]. These works show that minima for the same task are typically connected by simple low-loss curves, and that such connectivity benefits training for multi-task and continual learning. More specifically, the authors show that embracing such mode connectivity between models improves accuracy on each task and remedies the risk of catastrophic forgetting.

In this paper, we leverage such effects, and propose federated learning over connected modes (FLOCO), where the clients share and together train a *solution simplex*—a linearly connected low-loss region in the weight parameter space. Specifically, FLOCO represents clients as points within the standard simplex based on the similarity between their gradients, and assigns each client a specific subregion of the simplex. Clients then participate in FL by sampling different models within their assigned subregions and sending back the gradient information to update the vertices of the global solution simplex (see Fig.1). This method facilitates collaborative training through the common solution simplex, while allowing for client-specific personalization according to their local data distributions.

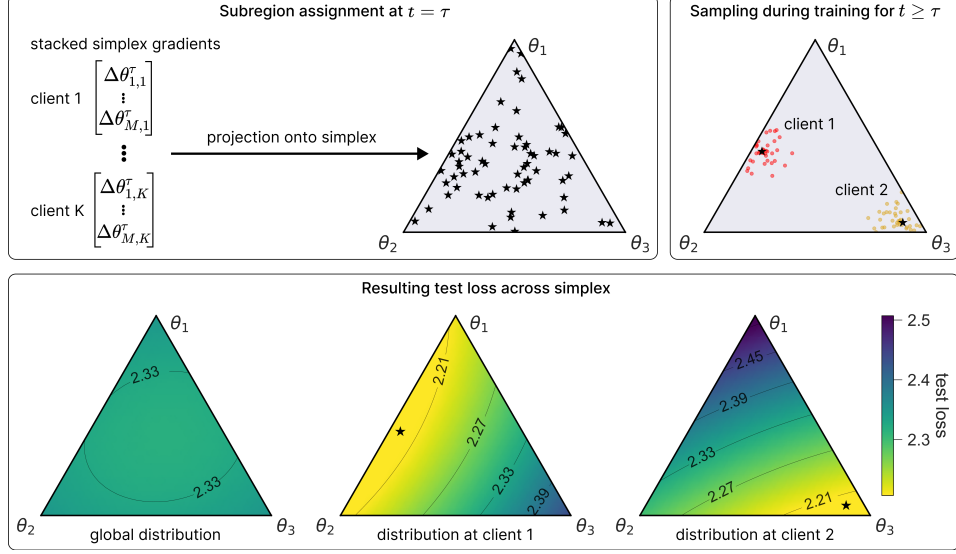


Figure 1: FLOCO expresses each client as a point (\star in the top-middle plot) by projecting the gradient signals onto the simplex, so that similar clients are close to each other. In each communication round, each client uniformly samples points in the neighborhood of their projected point (top-right plot), and jointly train the solution simplex. The lower row plots the resulting test loss on the solution simplex, where the loss for the global distribution (left) is uniformly small, while the losses for individual local distributions (middle for client 1 and right for client 2) are small around their projected points.

Our experiments show that FLOCO outperforms state-of-the-art approaches in both global (Fed-Prox [13]), and personalized (Ditto [14]) FL without introducing significant computational overhead. We also show other benefits of FLOCO, including better uncertainty estimation, improved worst client performance, and smaller divergence of the gradient signals.

Our main contributions are summarized as follows:

- We propose FLOCO, a novel FL method that trains a solution simplex for mitigating the statistical heterogeneity of clients, and demonstrate its state-of-the-art performance for global and local personalized FL.
- We propose a simple projection method to express clients as points in the standard simplex based on the gradient signals, and establish a procedure of subregion assignments.
- We conduct experimental evaluations on semi-artificial and real-world FL benchmarks with detailed analyses of the behavior of FLOCO, which give insights into the mechanism on how it improves performance compared to the baselines.

Our code is submitted as a supplement and will be made public upon acceptance.

2 Background

In this section, we briefly explain the concepts behind federated learning and mode connectivity, which form the backbone of our approach. The symbols that we use throughout the paper are listed in Table 5 in Appendix.

2.1 Federated Learning

Assume a federated system where the server has a global model g_0 and the K clients have their local models $\{g_k\}_{k=1}^K$. FL aims to obtain the best performing models $\{g_k^*\}_{k=0}^K$ such that

$$g_0^* = \operatorname{argmin}_{g_0} F^*(g_0) \equiv \sum_{k=1}^K p(k) F_k^*(g_0), \quad (1)$$

$$g_k^* = \operatorname{argmin}_{g_k} F_k^*(g_k) \text{ for } k = 1, \dots, K, \quad (2)$$

$$\text{where } F_k^*(g) = \mathbb{E}_{(\mathbf{x}, y) \sim p_k(\mathbf{x}, y)} [f(g, (\mathbf{x}, y))].$$

Here, $p(k)$ is the normalized population of data samples over the clients, $p_k(\mathbf{x}, y)$ is the data distribution for the client k , and $f(g, (\mathbf{x}, y))$ is the loss, e.g., cross-entropy, of the model g on a sample $(\mathbf{x}, y) \in \mathbb{R}^I \times \{1, \dots, L\}$, where I is the dimension of an input data sample. *Global* [15] and *personalized* [8] FL aim to approximate g_0^* and $\{g_k^*\}_{k=1}^K$, respectively, by using the training data $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^K$ observed by the clients. Throughout the paper, we assume that all models are neural networks (NNs) $\hat{y} = g_k(\mathbf{x}; \mathbf{w}_k)$ with the same architecture, and represent the model g_k with its NN weight parameters $\mathbf{w}_k \in \mathbb{R}^D$, i.e., we hereafter represent $g_k(\mathbf{x}; \mathbf{w}_k)$ by \mathbf{w}_k and thus denote, e.g., $F_k^*(g_k)$ by $F_k^*(\mathbf{w}_k)$. Let $N = \sum_{k=1}^K N_k$ be the total number of samples, where $N_k = |\mathcal{D}_k|$.

For the independent and identically distributed (IID) data setting, i.e., $p_k(\mathbf{x}, y) = p(\mathbf{x}, y), \forall k = 1, \dots, K$, the global and personalized FL aim for the same goal, and the minimum loss solution for the given training data is

$$\hat{\mathbf{w}}_0 = \hat{\mathbf{w}}_k = \operatorname{argmin}_{\mathbf{w}} F(\mathbf{w}) \equiv \sum_{k=1}^K \frac{N_k}{N} F_k(\mathbf{w}), \quad (3)$$

$$\text{where } F_k(\mathbf{w}) = \frac{1}{N_k} \sum_{(\mathbf{x}, y) \in \mathcal{D}_k} f(\mathbf{w}, (\mathbf{x}, y)).$$

In this setting, Federated Averaging (FedAvg) [1],

$$\mathbf{w}_0^{t+1} = \mathbf{w}_0^t + \sum_{k \in \mathcal{S}^t} \frac{N_k}{N} \cdot \Delta \mathbf{w}_k^{t+1} \text{ for } t = 1, \dots, T, \quad (4)$$

is known to converge to $\hat{\mathbf{w}}_0$, and thus solve Eq. (3). Here, \mathcal{S}^t is the set of clients that participate the t -th communication round, and $\Delta \mathbf{w}_k^{t+1} = \mathbf{w}_k^{t+1} - \mathbf{w}_k^t$ is the update after T' steps of the local gradient descent,

$$\check{\mathbf{w}}^{t'+1} = \check{\mathbf{w}}^{t'} - \gamma \nabla F_k(\check{\mathbf{w}}^{t'}), \text{ for } t' = 1, \dots, T', \quad (5)$$

where $\check{\mathbf{w}}^0 = \mathbf{w}_0^t$, $\check{\mathbf{w}}^{T'} = \mathbf{w}_k^{t+1}$, and γ is the step size. FedAvg has been further enhanced with, e.g., proximity regularization [13], auxiliary data [16], and ensembling [17].

On the other hand, in the more realistic non-IID setting, where $\mathbf{w}_0^* \neq \mathbf{w}_k^*$, FedAvg and its variants suffer from slow convergence and poor local performance [18]. To address such challenges, Ditto [14] was proposed for personalized FL, i.e., to approximate the best local models $\{\mathbf{w}_k^*\}_{k=1}^K$. Ditto has two training phases: it first trains the global model $\hat{\mathbf{w}}_0$ by FedAvg, then trains the local models with proximity regularization to $\hat{\mathbf{w}}_0$, i.e.,

$$\hat{\mathbf{w}}_k = \operatorname{argmin}_{\mathbf{w}_k} \tilde{F}_k(\mathbf{w}_k, \hat{\mathbf{w}}_0) \equiv F_k(\mathbf{w}_k) + \frac{\lambda}{2} \|\mathbf{w}_k - \hat{\mathbf{w}}_0\|_2^2,$$

where λ controls the divergence from the global model.

Ditto outperforms other non-IID FL methods, including the client clustering method HYPCLUSTER, adaptive federated learning (APFL), which interpolates between a global and local models [19], Loopless Local SGD (L2SGD), which applies global and local model average regularization [20], and MOCHA [6], which fits task-specific models through a multi-task objective. We refer readers to [8] for a comprehensive study of existing personalized FL methods.

2.2 Mode Connectivity and Solution Simplex

Freeman and Bruna (2017) [21], as well as Garipov et al. (2018) [10], discovered the mode connectivity in the NN parameter space—the existence of simple regions with low training loss between two well-trained models from different initializations. Nagarajan and Kolter (2019) [11] showed that the path is linear when the models are trained from the same initialization, but with different ordering of training data. Frankle et al. (2020) [22] showed that the same pre-trained models stay linearly-connected after fine-tuning with gradient noise or different data ordering.

Benton et al. (2021) [23] found that the low loss connection is not necessarily in 1D, and [24] showed that a simplex,

$$\mathcal{W}(\{\boldsymbol{\theta}_m\}) = \left\{ \mathbf{w}_\alpha(\{\boldsymbol{\theta}_m\}) = \sum_{m=1}^{M+1} \alpha_m \boldsymbol{\theta}_m; \boldsymbol{\alpha} \in \Delta^M \right\}, \quad (6)$$

within which any point has a small loss, can be trained from randomly initialized endpoints. Here, $\{\boldsymbol{\theta}_m \in \mathbb{R}^D\}_{m=1}^{M+1}$ are the endpoints or vertices of the simplex, and $\Delta^M = \{\boldsymbol{\alpha} \in [0, 1]^{M+1}; \|\boldsymbol{\alpha}\|_1 = 1\}$ denotes the M -dimensional standard simplex. This simplex learning is performed by finding the endpoints that (approximately) minimize

$$\mathbb{E}_{(\mathbf{x}, y) \sim p(\mathbf{x}, y)} [\mathbb{E}_{\mathbf{w} \sim \mathcal{U}_{\mathcal{W}(\{\boldsymbol{\theta}_m\})}} [f(\mathbf{w}, (\mathbf{x}, y))]], \quad (7)$$

where $\mathcal{U}_{\mathcal{W}}$ denotes the uniform distribution on a set \mathcal{W} .

3 Proposed Method

In this section, we introduce our approach, where the mode connectivity is leveraged for collaborative training between personalized client models.

3.1 Federated Learning over Connected Modes (FLOCO)

The main idea behind FLOCO is to assign subregions of the solution simplex (6) to clients in such a way that similar clients train neighboring (and overlapped) regions, while enforcing (linear) connectivity to all other client's subregions. The connectivity constraint systematically regularizes client training and allows for efficient collaboration between them.

The subregion assignments need to reflect the similarity between the clients. To this end, FLOCO expresses each client as a point in the standard simplex, based on the gradient update signals. Specifically, it applies the *Euclidean projection onto the positive simplex* [25] with the Riesz s-Energy regularization [26], which gives well spreaded projections that preserve the similarity between the client's gradient signals as much as possible. Once the clients are projected onto the standard simplex as $\{\alpha_k \in \Delta^M\}_{k=1}^K$, we assign the L1-ball with radius r around α_k , i.e., $\mathcal{R}_k = \{\alpha \in \Delta^M; \|\alpha - \alpha_k\|_1 \leq \rho\}$, to the k -th client. Note that the gradient update signals are informative for the subregion assignment only after the (global) model is trained to some extent. Therefore, the subregion assignment is performed after τ FL rounds are performed. Before the assignment, i.e., $t \leq \tau$, all clients train the whole standard simplex $\mathcal{R}_k = \Delta^M, \forall k$, which corresponds to a simplex learning version of FedAvg.

Starting from randomly initialized simplex endpoints $\{\theta_m\}_{m=1}^{M+1}$, FLOCO performs the following steps for each participating client $k \in \mathcal{S}^t$ in each communication round t :

1. The server sends the current endpoints $\{\theta_m^t\}_{m=1}^{M+1}$ to the client k .
2. The client k performs simplex learning only on the assigned subregion \mathcal{R}_k as a local update.
3. The client sends the local update of the endpoints to the server.

This way, FLOCO is expected to learn the global solution simplex $\{w_\alpha; \alpha \in \Delta^M\}$, while allowing personalization to local client distributions within the solution simplex. Algorithm 1 shows the main steps.

Although the simplex learning can be applied to all weight parameters, our preliminary experiment showed that simplex learning only of the weights in the last fully-connected layer (and the other weights are point-estimated) is sufficient. Therefore, our FLOCO only applies the simplex learning to the last layer, which gives other benefits including applicability to finetuning of pre-trained models, and significant reduction of computational and communication costs, see Section 4.3.

Below, we describe detailed procedures of client projection, local and global updates in the communication rounds, and inference in the test time.

3.2 Client Gradient Projection onto Standard Simplex

Here we explain how to obtain the representations $\{\alpha_k \in \Delta^M\}$ of the clients in the standard simplex such that similar clients are located close to each other, while all clients are well-spread across the simplex.

At communication round $t = \tau$, FLOCO uses the gradient updates of the endpoints $\{\Delta\theta_{m,k}^\tau\}_{m=1}^{M+1}$ as a representation of the client k . We concatenate the gradients for the $M + 1$ endpoints into a $((M + 1) \cdot D)$ -dimensional vector, and apply the PCA projection onto the M dimensional space, yielding $\kappa_k \in \mathbb{R}^M$ as a low dimensional representation. To project $\{\kappa_k\}$ onto the standard simplex Δ^M , we solve the following minimization problem:

$$\min_{z \geq 0} \sum_{i,j} \frac{1}{\|\beta_i(z) - \beta_j(z)\|_2^2}, \quad (8)$$

$$\text{subject to: } \hat{\beta}_k(z) = \operatorname{argmin}_{\beta_k \in \Delta^{M-1}} \|\beta_k - \kappa_k\|_2^2. \quad (9)$$

The objective function in Eq.(8) is the Riesz s-Energy [26], a generalization of potential energy of multiple particles in a physical space, and therefore its minimizer corresponds to the state where

Algorithm 1: Federated Learning over Connected Modes (FLOCO).

Input : number of communication rounds T , number of clients K , simplex dimension M ,
subregion assignment round τ , subregion radius ρ

```
1  $\{\theta_m^0\}_{m=1}^{M+1} \leftarrow \text{initialize\_simplex}(M)$ 
2  $\mathcal{R}_k \leftarrow \Delta^M, \forall k = 1, \dots, K$  // set all client subregions to the whole standard simplex
3 for  $t = 1$  to  $T$  do
4   if  $t = \tau$  then
5      $\{\{\Delta\theta_{m,k}^\tau\}_{m=1}^{M+1}\}_{k=1}^K \leftarrow \text{collect\_and\_stack\_gradients}()$ 
6      $\{\alpha_k\}_{k=1}^K \leftarrow \text{client\_representation}(\{\{\Delta\theta_{m,k}^\tau\}_{m=1}^{M+1}\}_{k=1}^K)$ 
7      $\{\mathcal{R}_k\}_{k=1}^K \leftarrow \text{assign\_subregions}(\{\alpha_k\}_{k=1}^K, \rho)$ 
8    $\mathcal{S}^t \leftarrow \text{choose\_participating\_clients}()$ 
9   for  $k \in \mathcal{S}^t$  do
10     $\{\theta_{m,k}^{t+1}\}_{m=1}^{M+1} \leftarrow \text{local\_update}(\{\theta_{m,k}^t\}_{m=1}^{M+1}, \mathcal{R}_k)$ 
11   $\{\theta_m^{t+1}\}_{m=1}^{M+1} \leftarrow \text{global\_update}(\{\{\theta_{m,k}^{t+1}\}_{m=1}^{M+1}\}_{k \in \mathcal{S}^t})$ 
```

particles are well spread across the space. The minimization in the constraint (9) corresponds to the *Euclidean projection onto the positive simplex* [25], which forces $\{\beta_k\}$ to keep the locations of the PCA projections $\{\kappa_k\}$ of the clients. Fortunately, this minimization problem (for a fixed z) is convex, and can be efficiently solved (see Appendix A). We solve the main problem (8) by computing $\hat{\beta}_k(z)$ on a 1D grid in $z \in [0, 1]$ with the interval 0.001, and set the representations of the clients to $\alpha_k = \frac{\hat{\beta}_k(\hat{z})}{\hat{z}}$, where \hat{z} is the minimizer of Eq.(8).

3.3 Communication Round: Local and Global Updates

In the t -th communication round, the server sends the current endpoints $\{\theta_m^t\}_{m=1}^{M+1}$ to the participating clients \mathcal{S}^t . Then, each client $k \in \mathcal{S}^t$ draws one sample per mini-batch from the uniform distribution $\mathcal{A} = \{\alpha_b\}_{b=1}^B \sim \mathcal{U}_{\mathcal{R}_k}$ on the assigned subregion,

and applies T' local updates,

$$\check{\theta}_m^{t'+1} = \check{\theta}_m^{t'} - \alpha_m \cdot \gamma \cdot \nabla F_k(w_\alpha), \quad (10)$$

to the endpoints with α sequentially chosen from \mathcal{A} .¹ Here $\check{\theta}_m^0 = \theta_m^t, \check{\theta}_m^{T'} = \theta_{m,k}^{t+1}$. The local updates $\{\Delta\theta_{m,k}^{t+1} = \theta_{m,k}^{t+1} - \theta_m^t\}_{m=1}^{M+1}$ are sent back to the server, which updates the endpoints as

$$\theta_m^{t+1} = \theta_m^t + \sum_{k \in \mathcal{S}^t} \frac{N_k}{N} \cdot \Delta\theta_{m,k}^{t+1}. \quad (11)$$

As explained in Section 3.1, the client subregions are initially set to the whole simplex Δ^M before the subregion assignment is performed at $t = \tau$, which corresponds to a straightforward application of the simplex learning to FedAvg. After the subregion assignment, FLOCO uses the degrees of freedom within the solution simplex to personalize clients models.

3.4 Inference

With the trained endpoints $\{\hat{\theta}_m = \theta_m^T\}_{m=1}^{M+1}$, we simply use $w_{\hat{\alpha}_0}(\{\hat{\theta}_m\}_{m=1}^{M+1})$ as the global model, where $\hat{\alpha}_0 = \frac{1}{M+1} \mathbf{1}_{M+1}$ with $\mathbf{1}_D$ denoting the D -dimensional all one vector. For local models, we use $\{w_{\hat{\alpha}_k}(\{\hat{\theta}_m\}_{m=1}^{M+1})\}_{k=1}^K$ where $\hat{\alpha}_k = \alpha_k$.

¹Note that we do not rely on any regularizer that forces the diversity of the endpoints, as in [24], because, in FLOCO, the diversity of local client distributions prevents the simplex endpoints from collapsing to a single point.

4 Experiments

In this section, we experimentally show the advantages of FLOCO over the baseline methods.

4.1 Experimental Setting

Datasets and models To evaluate FLOCO, we perform image classification on the CIFAR-10 [27] and the FEMNIST [28] datasets. For CIFAR-10, we train a simple CNN [29] from scratch, as well as fine-tune a ResNet-18 [30] pre-trained on ImageNet [31], as proposed in [32]. For FEMNIST, we train a simple CNN [29] from scratch as proposed in [29] as well as fine-tune a SqueezeNet [33] pre-trained on ImageNet as in [32]. We provide a table with the training hyperparameters that we use for each data set/model setting in Appendix B.

Data heterogeneity for non-FL benchmarks The FEMNIST dataset is an FL benchmark based on real data, where client heterogeneity is inherently embedded in the dataset. For CIFAR-10, we simulate statistical heterogeneity by two partitioning procedures. The first procedure by [34] partitions clients in equally sized groups and assigns each group a set of primary classes. Every client gets $z\%$ of its data from its group’s primary classes and $(100 - z)\%$ from the remaining classes. We apply this method with $z = 80$ for five groups and refer to this split as *5-Fold*. For example, in CIFAR-10 *5-Fold*, 20% of the clients get assigned 80% samples from classes 1-2 and 20% from classes 3-10. The second procedure, inspired by [35] and [36], draws the multinomial parameters of the client distributions $p_k(y) = \text{Multi}(y; \phi_k)$ from Dirichlet, i.e., $\phi_k \sim \text{Dir}_L(\beta)$, where β is the concentration parameter controlling the sparsity and heterogeneity— $\beta \rightarrow \infty$ concentrates the mass to the uniform distribution (and thus homogeneous), while small $0 < \beta < 1$ generates sparse and heterogeneous non-IID client distributions.

Baseline methods Besides FedAvg [1], we chose the most important and state-of-the-art methods for global and personalized FL as baselines. For global FL methods we chose FedProx [13], Feddyn and FedBla, while for personalized FL methods we chose Ditto [14], APFL, pFedSim and FedPAC.

FLOCO Hyperparameters For simple CNN on the simulated non-IID splits Dir(0.5) and Five-Fold, we set $\tau = 100$, $M = 6$, $\rho = 0.1$. For simple CNN on FEMNIST we set $\tau = 200$, $M = 3$, $\rho = 0.3$. For pre-trained ResNet-18 on the simulated non-IID splits Dir(0.5) and Five-Fold we set $\tau = 50$, $M = 6$, $\rho = 0.1$ and for the pre-trained SqueezeNet on FEMNIST we set $\tau = 20$, $M = 3$, $\rho = 0.05$. We found those settings work well in our preliminary experiments, and conducted ablation study with other parameter settings in Appendix D. For the baselines, we follow the recommended parameter settings by the authors, which are detailed in Appendix B.

Evaluation criteria For the performance evaluation, we adopt two metrics, the test accuracy measured after the last communication round (ACC) and the time-to-best-accuracy (TTA), each for evaluating the global and local FL performance. ACC is the last test accuracy over T communication rounds, i.e., $\text{ACC}(T) = \frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \mathbb{1}(y_i = \arg\max g(\mathbf{x}_i; \hat{\mathbf{w}}^T))$, where $\mathbb{1}(\cdot)$ is the indicator function that equals to 1 if the event is true and 0 otherwise. TTA evaluates the number of communication rounds needed to achieve the best baseline (FedAvg and Ditto in this paper) test accuracy, i.e., $\text{ACC}_{\text{FedAvg}}(T)$. We report TTA improvement, i.e. the TTA of the baseline, e.g. FedAvg, divided by the TTA of the benchmarked method, e.g. FLOCO. Moreover, we report the expected-calibration-error (ECE) [37], a common measure that evaluates the quality of uncertainty estimation of a trained model, for the last communication round.

4.2 Results

Table 1 and 2 summarize the main experimental results, where FLOCO consistently outperforms the baselines across the different experiments in terms of global (red) and local (blue) test accuracy, as well as test ECE. The global and local test metrics are measured after the last communication round and averaged over 5 different seed runs. Below we report on detailed observations.

Global and local FL test accuracy. We first evaluate the global and local test performance on CIFAR-10 with the non-IID data splits generated by the 5-Fold and Dir(β) procedures, as well

Table 1: Average **global** and **local** test accuracy.

	CIFAR-10								FEMNIST			
	SimpleCNN				pre-trained ResNet-18				SimpleCNN		pre-trained SqueezeNet	
	5-Fold		Dir(0.5)		5-Fold		Dir(0.5)					
FedAvg	66.65	67.08	68.22	68.74	67.35	67.74	76.17	76.49	77.47	79.75	76.80	78.79
FedProx	66.14	66.80	67.79	67.67	67.88	68.00	75.72	75.88	76.91	78.21	76.39	78.06
Ditto	66.77	79.69	67.96	74.46	69.01	83.29	75.84	80.53	76.92	80.05	76.90	83.12
FLOCO	68.26	80.92	69.79	74.64	74.61	87.38	79.11	82.29	77.95	83.89	77.14	83.54

Table 2: Average **global** and **local** expected test calibration error.

	CIFAR-10								FEMNIST			
	SimpleCNN				pre-trained ResNet-18				SimpleCNN		pre-trained SqueezeNet	
	5-Fold		Dir(0.5)		5-Fold		Dir(0.5)					
FedAvg	14.73	17.03	14.08	16.61	5.50	13.02	2.12	10.20	14.02	17.11	13.80	17.33
FedProx	14.73	16.48	13.64	16.67	5.00	12.60	1.93	10.19	13.47	17.70	13.39	17.71
Ditto	14.68	11.56	14.23	14.18	5.00	9.73	1.96	9.66	14.20	16.38	13.74	13.93
FLOCO	10.98	9.84	10.12	12.95	2.68	7.62	1.91	9.21	12.96	13.95	13.28	13.33

Table 3: Average **local** test accuracy for the 5 % worst performing clients.

	CIFAR-10				FEMNIST	
	SimpleCNN		pre-trained ResNet-18		SimpleCNN	pre-trained SqueezeNet
	5-Fold	Dir(0.5)	5-Fold	Dir(0.5)		
FedAvg	0.42 ± 0.03	0.51 ± 0.04	0.50 ± 0.06	0.65 ± 0.03	0.46 ± 0.08	0.42 ± 0.12
FedProx	0.45 ± 0.05	0.51 ± 0.04	0.52 ± 0.02	0.64 ± 0.03	0.44 ± 0.10	0.42 ± 0.09
Ditto	0.65 ± 0.02	0.60 ± 0.03	0.72 ± 0.04	0.67 ± 0.05	0.37 ± 0.20	0.45 ± 0.19
FLOCO	0.66 ± 0.03	0.59 ± 0.03	0.79 ± 0.02	0.72 ± 0.03	0.55 ± 0.07	0.53 ± 0.13

as the natural non-IID data splits in the FEMNIST dataset. Table 1 shows the test accuracies on CIFAR-10 with SimpleCNN trained from random initialization (left) and ResNet-18 fine-tuned from the ImageNet pre-trained model (middle), respectively. It also shows the test accuracies on FEMNIST with SimpleCNN trained from random initialization and SqueezeNet fine-tuned from the ImageNet pre-trained model (right). We clearly see that FLOCO outperforms all baselines in terms of global (red) as well as average local (blue) test accuracy.

In some cases, e.g., SimpleCNN on the CIFAR-10 Dir(0.5) split, as well as pre-trained SqueezeNet on the FEMNIST, the performance gain from the DITTO baseline is marginal. However, FLOCO is still beneficial, because training DITTO requires roughly 2 times larger computation and storage costs than FLOCO, since the overhead added to FEDAVG by FLOCO is minimal, see Section 4.3. Furthermore, the personalization procedure of DITTO can be straightforwardly integrated into the FLOCO framework, boosting its performance and significantly outperforming DITTO (We call this method FLOCO+, which is introduced and evaluated in Appendix C). Besides test accuracy, we also observe TTA improvement of our method for all settings of up to $\times 4$ for average local test accuracy and $\times 3.8$ for global test accuracy than the next best baselines. We report all TTAs in Table 4.

Calibration. Next, we evaluate and benchmark the quality of uncertainty estimation of all methods. For this purpose we evaluate the ECE on each model-dataset combination for each baselines and show the results in Table 2. As shown, FLOCO yields better ECE for every setting, ranging from 1.23-3.52 and 0.05-3.52 of local and global ECE improvement compared to the next best baseline, respectively.

Worst client performance. Lastly, we evaluate the average local, as well as global test accuracies of the worst 5% of clients, as is commonly done in order to measure the biasedness of the FL method towards certain clients or client groups [38]. We evaluate the worst 5% client performance on all model-dataset combinations over 5 trial runs and display the results in Table 3. We observe that FLOCO yields the best worst client performance across all benchmarks ranging from 1-9% percent

improvement compared to the next best baseline, except for SimpleCNN on CIFAR-10 with the Dir(0.5) split, where FLOCO very closely ($\sim 1\%$ difference) approximates the performance of DITTO.

Time-to-accuracy. Similar to Table 1, we plot the TTA improvement for FLOCO. In particular, we show the TTA improvement of FLOCO over FedAvg and Ditto. Again, all reported results are average values that we computed over 5 different random seeds.

Table 4: Improvements for **global** and **local** time-to-accuracy.

	CIFAR-10								FEMNIST			
	SimpleCNN				pre-trained ResNet-18				SimpleCNN		pre-trained SqueezeNet	
	5-Fold		Dir(0.5)		5-Fold		Dir(0.5)					
FLOCO vs. FedAvg	x1.5	x4.4	x1.6	x3.7	x3.5	x4.33	x1.8	x2.6	x1.4	x2	x1.1	x2.4
FLOCO vs. Ditto	x1.5	x1.2	x1.3	x1.8	x3.8	x1.9	x2.7	x1.9	x2.4	x1.9	x1	x1.1

4.3 Analysis and Discussion

Here we provide additional analyses and discussion on FLOCO.

Solution structure in simplex. First, we confirm that FLOCO uses the degrees of freedom within the solution simplex for personalization. To this end, we draw approximately 500 uniformly distributed points in the solution simplex, and evaluate the global and the local test accuracy of the corresponding models. Figure 1 (bottom row) shows the global test accuracy (left most) and the local test accuracy (center and right) for two clients. As expected, for the global test dataset the solution simplex performs uniformly well across all its area, while the losses for the two individual local client distributions are small around their projected points (\star). This result indicates that the heterogeneous sharing of the solution simplex across the clients properly works as designed.

Gradient variance reduction and stability of training. Figure 2 shows the test accuracy curves during training for global (left) and local (center) test accuracies of different methods with the standard deviation over 5 trials as shadows. We observe that FLOCO not only converges to a better global and local test accuracy, but also shows small standard deviation across trials. The latter implies that our systematic regularization through the solution simplex stabilizes the training dynamics significantly. Figure 2 (right) shows the total gradient variance—the sum of the variances of the updates $\Delta \mathbf{w}_k^t = \mathbf{w}_k^t - \mathbf{w}_0^{t-1}$ for FedAvg, and $\Delta \bar{\boldsymbol{\theta}}_{m,k}^t = \bar{\boldsymbol{\theta}}_{m,k}^t - \bar{\boldsymbol{\theta}}_{m,0}^{t-1}$ for FLOCO, respectively. More specifically, we compute the variance over the last fully-connected layer, given by

$$\sigma_{\text{total}}^2(t) = \sum_{k \in \mathcal{S}^t} \|\Delta \mathbf{w}_k^t - \frac{1}{|\mathcal{S}^t|} \sum_{k \in \mathcal{S}^t} \Delta \mathbf{w}_k^t\|_2^2, \quad (12)$$

for FedAvg and FedProx, and by

$$\sigma_{\text{total}}^2(t) = \sum_{k \in \mathcal{S}^t} \|\Delta \bar{\boldsymbol{\theta}}_k^t - \frac{1}{|\mathcal{S}^t|} \sum_{k \in \mathcal{S}^t} \Delta \bar{\boldsymbol{\theta}}_k^t\|_2^2, \quad (13)$$

with $\Delta \bar{\boldsymbol{\theta}}_k^t = \frac{1}{M+1} \sum_{m=1}^{M+1} \Delta \bar{\boldsymbol{\theta}}_{m,k}^t$, for FLOCO. As discussed in [39, 40], a small total variance indicates effective collaborations with consistent gradient signals between the clients, leading to better performance. From the figure, we see that the total gradient variance of FLOCO is much lower and more stable than the baseline methods, which, together with its good performance observed in Table 1, is consistent with their discussion. The variance reduction with FLOCO implies that the degrees of freedom of the solution simplex can absorb the heterogeneity of clients to some extent, making the gradient signals more homogeneous. Moreover, [41] argued that the last classification layer has the biggest impact on performance, implying that reducing the total variance of the classification layer, as FLOCO does with simplex learning, is most effective.

Computational complexity. If the batch size is one, simplex training adds $O(\pi \cdot M)$ computational complexity for each layer, where π is the parameter complexity of the layer, e.g., $\pi = d \cdot L$ for a fully connected layer with d input and L output neurons, and M is the simplex dimension [24]. For FLOCO, this additional complexity only applies to the classification layer. For inference, no additional complexity arises, compared to FedAvg, because inference is performed by the single




Figure 2: Global (left) and average local (center) test accuracy for SimpleCNN on FEMNIST. We can clearly observe a jump in average local test accuracy at $\tau = 200$, which is a result of our subregion assignment. Right shows the total variance of the gradients for the last fully-connected layer.

model corresponding to the cluster center. Since the most modern architectures, e.g., ResNet-18 and Vision Transformer (ViT) [42], have parameter complexity of $O(\mathbb{G}_{\text{FE}}) \gg O(\mathbb{G}_{\text{C}})$, where \mathbb{G}_{FE} and \mathbb{G}_{C} are the complexities of the feature extractor and the classification layer, respectively, the additional training complexity, applied only to the classification layer, of FLOCO is ignorable, i.e., $O(\mathbb{G}_{\text{FE}}) \gg O(\mathbb{G}_{\text{C}} \cdot M)$. The same applies to the communication costs: since the simplex learning is applied only to the classification layer, the increase of communication costs are ignorable compared to the communication costs for the feature extractor.

5 Related Work

There are a few existing works that apply simplex learning to federated learning. [29] showed that enforcing a low loss simplex between independently initialized global and client models yields good personalized FL performance. Their approach builds on [19], which finds optimal interpolation coefficients between a global and local model to improve personalized FL. However, their simplex is restricted to be 1D, i.e., a line segment, and the global model performance is comparable to the plain FedAvg. Moreover, they train a solution simplex over all layers between global and local models, which is computationally expensive and limits its applicability to training *from scratch*. This should be avoided if pre-trained models are available [32, 43]. Our method generalizes to training low-loss simplices of higher dimensions in a FL setting, tackles both the global and personalized FL objectives, is applicable to pre-trained models, and shows significant performance gains by employing our proposed subregion assignment procedure.

6 Conclusion

FL on highly non-IID client data distributions remains a challenging problem and a very actively researched topic. Recent works tackle non-IID FL settings either through global or personalized FL. While the former aims to find a single optimal set of parameters that fit a global objective, the latter tries to optimize multiple local models each of which fits the local distributions well. These two different objectives may pose a trade-off, that is, personalized FL might adapt models to strongly to local distributions which might harm the global performance, while global FL solutions might fit none of the local distributions if the local distributions are diverse. In this paper, we addressed this issue by leveraging the mode-connectivity of neural networks. Specifically, we propose FLOCO, where each client trains an assigned subregion within the solution simplex, which allows for personalization, and at the same, contributes to learning a well-performing global model. FLOCO achieves state-of-the-art performance in both global and personalized FL, with minimal computational and communication overhead during training and no overhead during inference.

Promising future research directions include better understanding the decision-making process of solution simplex training through global and local explainable AI methods [44–46]. Furthermore, we want to apply our approach to continual learning problems and FL scenarios with highly varying client availability [47, 48].

References

- [1] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282, 2017.
- [2] Viraj Kulkarni, Milind Kulkarni, and Aniruddha Pant. Survey of personalization techniques for federated learning. In *2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*, pages 794–797. IEEE, 2020.
- [3] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3710–3722, 2020.
- [4] Fan Lai, Xiangfeng Zhu, Harsha V Madhyastha, and Mosharaf Chowdhury. Oort: Efficient federated learning via guided participant selection. In *Symposium on Operating Systems Design and Implementation*, pages 19–35, 2021.
- [5] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022.
- [6] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. *Advances in Neural Information Processing Systems*, 30, 2017.
- [7] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- [8] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [9] Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *International Conference on Machine Learning*, pages 1309–1318, 2018.
- [10] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikin, Dmitry P Vetrov, and Andrew G Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. *Advances in Neural Information Processing Systems*, 31, 2018.
- [11] Vaishnavh Nagarajan and J Zico Kolter. Uniform convergence may be unable to explain generalization in deep learning. *Advances in Neural Information Processing Systems*, 32, 2019.
- [12] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Dilan Gorur, Razvan Pascanu, and Hassan Ghasemzadeh. Linear mode connectivity in multitask and continual learning. In *International Conference on Learning Representations*, 2021.
- [13] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Machine Learning and Systems*, 2:429–450, 2020.
- [14] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pages 6357–6368, 2021.
- [15] Chen Zhang, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.
- [16] Felix Sattler, Tim Korjakow, Roman Rischke, and Wojciech Samek. Fedaux: Leveraging unlabeled auxiliary data in federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(9):5531–5543, 2023.
- [17] Naichen Shi, Fan Lai, Raed Al Kontar, and Mosharaf Chowdhury. Fed-ensemble: Ensemble models in federated learning for improved generalization and uncertainty quantification. *IEEE Transactions on Automation Science and Engineering*, 2023.

- [18] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021.
- [19] Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*, 2020.
- [20] Filip Hanzely and Peter Richtárik. Federated learning of a mixture of global and local models. *arXiv preprint arXiv:1812.01097*, 2020.
- [21] C. Daniel Freeman and Joan Bruna. Topology and geometry of half-rectified network optimization. In *International Conference on Learning Representations*, 2017.
- [22] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pages 3259–3269, 2020.
- [23] Gregory Benton, Wesley Maddox, Sanae Lotfi, and Andrew Gordon Gordon Wilson. Loss surface simplexes for mode connecting volumes and fast ensembling. In *International Conference on Machine Learning*, pages 769–779, 2021.
- [24] Mitchell Wortsman, Maxwell C Horton, Carlos Guestrin, Ali Farhadi, and Mohammad Rastegari. Learning neural network subspaces. In *International Conference on Machine Learning*, pages 11217–11227, 2021.
- [25] Mathieu Blondel, Akinori Fujino, and Naonori Ueda. Large-scale multiclass support vector machine training via euclidean projection onto the simplex. In *2014 22nd International Conference on Pattern Recognition*, pages 1289–1294. IEEE, 2014.
- [26] Douglas P Hardin and Edward B Saff. Minimal riesz energy point configurations for rectifiable d-dimensional manifolds. *Advances in Mathematics*, 193(1):174–204, 2005.
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [28] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*, 2018.
- [29] Seok-Ju Hahn, Minwoo Jeong, and Junghye Lee. Connecting low-loss subspace for personalized federated learning. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 505–515, 2022.
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [31] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 248–255. Ieee, 2009.
- [32] John Nguyen, Jianyu Wang, Kshitiz Malik, Maziar Sanjabi, and Michael Rabbat. Where to begin? on the impact of pre-training and initialization in federated learning. In *International Conference on Learning Representations*, 2023.
- [33] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [34] Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized cross-silo federated learning on non-iid data. In *AAAI Conference on Artificial Intelligence*, volume 35, pages 7865–7873, 2021.
- [35] Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *International Conference on Machine Learning*, pages 7252–7261, 2019.

- [36] Liang Gao, Huazhu Fu, Li Li, Yingwen Chen, Ming Xu, and Cheng-Zhong Xu. Feddc: Federated learning with non-iid data via local drift decoupling and correction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [37] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330, 2017.
- [38] Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. In *International Conference on Learning Representations*, 2020.
- [39] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021.
- [40] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143, 2020.
- [41] Bo Li, Mikkel N Schmidt, Tommy S Alstrøm, and Sebastian U Stich. On the effectiveness of partial variance reduction in federated learning with heterogeneous data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3964–3973, 2023.
- [42] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021.
- [43] Hong-You Chen, Cheng-Hao Tu, Ziwei Li, Han Wei Shen, and Wei-Lun Chao. On the importance and applicability of pre-training for federated learning. In *International Conference on Learning Representations*, 2022.
- [44] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLOS ONE*, 10(7):1–46, 07 2015.
- [45] Wojciech Samek, Grégoire Montavon, Sebastian Lapuschkin, Christopher J. Anders, and Klaus-Robert Müller. Explaining deep neural networks and beyond: A review of methods and applications. *Proceedings of the IEEE*, 109(3):247–278, 2021.
- [46] Kirill Bykov, Mayukh Deb, Dennis Grinwald, Klaus-Robert Müller, and Marina M-C Höhne. Dora: Exploring outlier representations in deep neural networks. *Transactions on Machine Learning Research*, 2023.
- [47] A. Rodio, F. Faticanti, O. Marfoq, G. Neglia, and E. Leonardi. Federated learning under heterogeneous and correlated client availability. In *IEEE International Conference on Computer Communications*, 2023.
- [48] Philipp Wiesner, Ramin Khalili, Dennis Grinwald, Pratik Agrawal, Lauritz Thamsen, and Odej Kao. Fedzero: Leveraging renewable excess energy in federated learning. In *International Conference on Future and Sustainable Energy Systems*. ACM, 2024.
- [49] Richard L Burden and J Douglas Faires. 2.1 the bisection algorithm. *Numerical analysis*, 3, 1985.

Appendix

This appendix provides a nomenclature, details to our optimization problem and experimental setup, as well as additional results and insights.

Table 5: Nomenclature.

Symbol	Description
$k = 1, \dots, K$	Clients
$t = 1, \dots, T$	Communication rounds
$t' = 1, \dots, T'$	Local training iterations
S^t	Participating clients in round t
B	Mini-batch size
γ	Client gradient descent step size
\mathcal{D}_k	Training data of client k
N	Total number of samples
N_k	Number of samples at client k
$p_k(\mathbf{x}, y)$	data distribution of client k
$\mathbf{w}_0^t \in \mathbb{R}^D$	Global model at round t
$\mathbf{w}_k^t \in \mathbb{R}^D$	Model of client k at round t
$\Delta^M = \{\boldsymbol{\alpha} \in [0, 1]^{M+1}; \ \boldsymbol{\alpha}\ _1 = 1\}$	M -dimensional standard simplex
$\boldsymbol{\theta}_1^t, \dots, \boldsymbol{\theta}_{M+1}^t$	Simplex endpoints at round t
$\mathbf{w}_\alpha = \sum_{m=1}^M \alpha_m \boldsymbol{\theta}_m$	Model parameters at a point $\boldsymbol{\alpha} \in \Delta^M$
ρ	Subregion radius
\mathcal{R}_k	Assigned subregion of client k
$\tau \in [1, \dots, T]$	Subregion assignment round
$\boldsymbol{\kappa}_k \in \mathbb{R}^M$	Low dimensional representation of stacked gradient update $\{\Delta \boldsymbol{\theta}_{m,k}^\tau\}_{m=1}^{M+1}$ of client k

A Optimization Problem

The Lagrangian of the lower-level optimization problem in (9) has the following formulation $\mathcal{L}(\boldsymbol{\alpha}_k, \lambda) = \frac{1}{2} \|\boldsymbol{\alpha}_k - \boldsymbol{\kappa}_k\|_2^2 + \lambda(\mathbf{1}^T \boldsymbol{\alpha}_k - z)$, with $\lambda \in \mathbb{R}$ being the Langrange multiplier. The Lagrangian can be further rewritten to $\mathcal{L}(\boldsymbol{\alpha}_k, \lambda) = \frac{1}{2} \|\boldsymbol{\alpha}_k - (\boldsymbol{\kappa}_k - \lambda \mathbf{1})\|_2^2 + \lambda(\mathbf{1}^T \boldsymbol{\kappa}_k - z) - \lambda^2 n$ such that the optimization problem reduces to solving:

$$\min_{z \in \mathbb{R}} \frac{1}{2} \|\boldsymbol{\alpha}_k - (\boldsymbol{\kappa}_k - \lambda \mathbf{1})\|_2^2 \quad (14)$$

$$\text{subject to: } \boldsymbol{\alpha}_k \succeq \mathbf{0}. \quad (15)$$

The optimal solution of (14) is given by $\boldsymbol{\alpha}_k^* = [\boldsymbol{\kappa}_k - \lambda^* \mathbf{1}]_+$. Plugging it back into the Lagrangian we get the following dual function

$$\mathcal{L}(\boldsymbol{\alpha}_k, \lambda) = \frac{1}{2} \|[\boldsymbol{\kappa}_k - \lambda^* \mathbf{1}]_+ - (\boldsymbol{\kappa}_k - \lambda \mathbf{1})\|_2^2 + \lambda(\mathbf{1}^T \boldsymbol{\kappa}_k - z) - \lambda^2 n \quad (16)$$

$$= \frac{1}{2} \|[\boldsymbol{\kappa}_k - \lambda^* \mathbf{1}]_+ - \boldsymbol{\kappa}_k + \lambda \mathbf{1}\|_2^2 + \lambda(\mathbf{1}^T \boldsymbol{\kappa}_k - z) - \lambda^2 n. \quad (17)$$

Finding $\boldsymbol{\alpha}_k^*$ can be achieved by maximizing (17) using for example the bisection algorithm [49]. After that the projected points are obtained as $\boldsymbol{\alpha}_k^* = [\boldsymbol{\kappa}_k - \lambda^* \mathbf{1}]_+$.

B Training Hyperparameters

Table 6 summarizes all hyperparameters that were used for each dataset/model combination. We train all models for a total of 500 communication rounds, except for ResNet-18 on CIFAR-10, which we train for 100 communication rounds. Moreover, we train each setting using a total of 100 clients, and for FEMNIST we select a randomly chosen subset of 100 total clients for each trial, of which we select 10 randomly to participate in training in each communication round, except for SimpleCNN on CIFAR-10 where we select 30 out of 100 clients to participate in each round. We evaluate all clients after every ten communication rounds. For CIFAR-10 we train a SimpleCNN with batch size 50 using SGD with a learning rate of 0.02, momentum of 0.5, and weight decay of 10^{-5} , and a pre-trained ResNet-18 with learning rate of batch size 32, using SGD with a learning rate of 0.01,

momentum of 0.9, and weight decay of 10^{-4} . For FEMNIST we train a pre-trained SqueezeNet with batch size 32 using SGD with a learning rate of 0.001, momentum of 0.9, weight decay of 10^{-4} , and a SimpleCNN with batch size 32, learning rate 0.1, momentum of 0, weight decay of 0. For FedProx we set the proximity hyperparameter to $\mu = 0.01$ for SimpleCNN on CIFAR-10, $\mu = 0.1$ for pre-trained ResNet-18 on CIFAR-10, SimpleCNN on FEMNIST, and pre-trained SqueezeNet on FEMNIST. For DITTO we set the local epochs to the same value as epochs for the global model $E_{\text{DITTO}} = E$. All training hyperparameters for CIFAR-10 and FEMNIST on a SimpleCNN were taken from [29], CIFAR-10 on a pre-trained ResNet-18 from and FEMNIST on pre-trained SqueezeNet from [32].

Table 6: Summary of used hyperparameters for training.

Dataset/Model	T	K	$ S^t $	e	E/E_{DITTO}	γ	mom.	wd	μ
CIFAR-10/SimpleCNN	500	100	30	50	5	0.02	0.5	10^{-5}	0.01
CIFAR-10/ResNet-18	100	100	10	32	5	0.01	0.9	10^{-4}	0.1
FEMNIST/SimpleCNN	500	100	10	32	5	0.1	0.0	0.0	0.1
FEMNIST/SqueezeNet	500	100	10	32	5	0.001	0.9	10^{-4}	0.1

C Combining FLOCO and Ditto

We can further enhance the personalized FL performance of FLOCO by additionally training a local model as in Ditto [14]. In this extension, called FLOCO⁺, each client personalizes the global endpoints $\{\hat{\theta}_m^0 = \theta_m\}_{m=1}^M$ by local gradient descent to minimize the Ditto objective, i.e.,

$$\begin{aligned} \{\hat{\theta}_m^k\} &= \operatorname{argmin}_{\{\theta_m\}} \tilde{F}_k(\{\theta_m\}, \{\hat{\theta}_m^0\}) \\ &\equiv \mathbb{E}_{\alpha \sim \mathcal{U}_{\mathcal{R}_{z_k}}} [F_k(w_\alpha(\{\theta_m\}))] + \frac{\lambda}{2} \sum_{m=1}^{M+1} \|\theta_m - \hat{\theta}_m^0\|_2^2. \end{aligned}$$

In Figure 3, we present preliminary results for FLOCO⁺ using the same hyperparameters as for FLOCO ($M = 6$, $\tau = 100$, and $\rho = 0.1$). We evaluate FLOCO⁺ for SimpleCNN on CIFAR-10 with Dir(0.5) split because, in this setting, FLOCO only slightly surpasses the average local accuracy of Ditto in our main results (Table 1). We observe that FLOCO⁺ outperforms both FLOCO and Ditto. Note that this improvement comes with additional computational cost.




Figure 3: Average local test accuracy for FLOCO⁺, FLOCO and baselines when training SimpleCNN on CIFAR-10 with Dir(0.5) split.

D Sensitivity to Parameter Setting

We investigate how stable the performance of FLOCO is for different hyperparameter settings. Specifically, we tested FLOCO with the combination of $\tau = 50, 100, 200$ (subregion assignment time step) and $\rho = 0.1, 0.2, 0.4$ (radius of subregions), and show the average local client and global test accuracy for SimpleCNN on CIFAR-10 5-Fold in Figure 4. We observe that the average local

client test accuracy (left) increases for earlier subspace projection starting points τ and lower client subregion radiuses ρ , with the best reached test accuracy being approximately 4% better than the worst, i.e., 82.79% against 78.18%. The intuition for this is that earlier client specialization in less overlapping regions allows for better personalization. On the other hand, as can be observed in the right heatmap of Figure 4 the global test performance is less sensitive to the choice of these hyperparameters, i.e., 70.66% against 69.30%. This is because the solution simplex is strained uniformly at every position even after projection trained, and thus the simplex midpoint is less sensitive to the specialization procedure to client distributions.




Figure 4: Local average client (left) and global (right) test accuracies for different subregion assignment time step τ and subregion radius ρ settings.

E Table 1: Flower Framework

E.1 Global test accuracy

Model-Dataset-Split	FLOCO	SuPerFed
SimpleCNN - CIFAR10 - 5-Fold	68.26	63.22
SimpleCNN - CIFAR10 - Dir(0.5)	69.79	63.00
SimpleCNN - FEMNIST - -	77.95	76.80
PretrainedResnet18 - CIFAR10 - 5-Fold	74.61	64.88
PretrainedResnet18 - CIFAR10 - Dir(0.5)	79.11	76.04
PretrainedSqueezeNet - FEMNIST - -	-	-

E.2 Average local test accuracy.

Model-Dataset-Split	FLOCO	SuPerFed
SimpleCNN - CIFAR10 - 5-Fold	80.92	75.65
SimpleCNN - CIFAR10 - Dir(0.5)	74.64	71.73
SimpleCNN - FEMNIST - -	83.89	82.59
PretrainedResnet18 - CIFAR10 - 5-Fold	87.38	52.78
PretrainedResnet18 - CIFAR10 - Dir(0.5)	82.29	60.91
PretrainedSqueezeNet - FEMNIST - -	-	-

F Table 2: FL-bench framework

F.1 Global test accuracy

Model-Dataset-Split	FLOCO	Ditto	FedAvg	FedRep	FedPac	FedDyn
SimpleCNN - CIFAR10 - Dir(0.5)	73.44	73.06	72.94	43.70	68.70	31.67

F.2 Average local test accuracy.

Model-Dataset-Split	FLOCO	Ditto	FedAvg	FedRep	FedPac	FedDyn
SimpleCNN - CIFAR10 - Dir(0.5)	79.54	79.20	73.42	68.00	73.15	31.39