

# FedDr+: Stabilizing Dot-regression with Global Feature Distillation for Federated Learning

Anonymous authors

Paper under double-blind review

## Abstract

Federated Learning (FL) has emerged as a pivotal framework for the development of effective global models (global FL) or personalized models (personalized FL) across clients with heterogeneous, non-iid data distribution. A key challenge in FL is client drift, where data heterogeneity impedes the aggregation of scattered knowledge. Recent studies have tackled the client drift issue by identifying significant divergence in the last linear (classifier) layer. To mitigate this divergence, strategies such as freezing the classifier weights and aligning the feature extractor accordingly have proven effective. Although the local alignment between classifier and feature extractor has been studied as a crucial factor in FL, we observe that it may lead the model to overemphasize the observed classes and underestimate the unobserved classes within each client. Therefore, our goals are twofold: (1) *improving local alignment* and (2) *maintaining the representation of unseen class samples*, ensuring that the solution seamlessly incorporates knowledge from individual clients, thus enhancing performance in both global and personalized FL. To achieve this, we introduce a novel algorithm named **FedDr+**, which empowers local model alignment using dot-regression loss. **FedDr+** freezes the classifier as a simplex ETF to align the features and improves aggregated global models by employing a feature distillation mechanism to retain information about unseen/missing classes. Our empirical results demonstrate that **FedDr+** not only outperforms methods with a frozen classifier but also surpasses other state-of-the-art approaches, ensuring robust performance across diverse data distributions.

## 1 Introduction

Federated Learning (FL) (McMahan et al., 2017; He et al., 2020b) is a distributed learning strategy that enables multiple clients to collaboratively train a model while preserving data privacy. The foundational method, FedAvg (McMahan et al., 2017), involves distributing a global model, training local models on each client’s private data, and aggregating these models without transmitting raw data. However, a major challenge in FL is data heterogeneity, or *non-iidness*, where differing data distributions across clients lead to *client drift*, hindering the convergence and effectiveness of the global model. Addressing this challenge requires focusing on two critical aspects: *local alignment* and *global knowledge preservation*. Local alignment ensures that the features extracted by the local model are properly aligned with the true class vectors of the classifier, enabling effective learning from the local data, while global knowledge preservation maintains the shared knowledge across clients, especially for classes unobserved in local training.

While both local alignment and global knowledge preservation are essential, they have generally been studied separately. Global knowledge preservation is crucial because it prevents the model from becoming overly biased toward the data of individual clients, ensuring decisions are based on a broader, shared understanding. This approach enables better generalization across all clients, particularly for unseen classes (Lee et al., 2022; 2024). The challenge of balancing global and local knowledge in FL resembles Catastrophic Forgetting in Continual Learning (CL) (McCloskey & Cohen, 1989), where learning new tasks can cause models to forget previously learned ones. To achieve this, strategies like FedProx (Li et al., 2020), MOON (Li et al., 2021a), and FedNTD (Lee et al., 2022) integrate global model regularization during local training. These methods

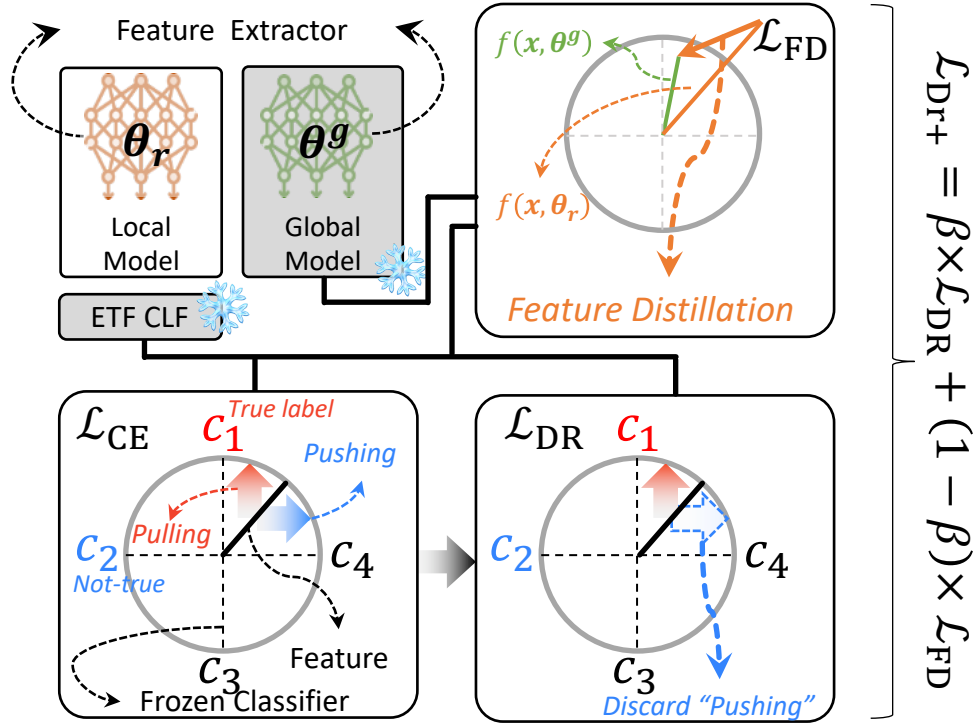


Figure 1: Overview of the proposed method, **FedDr+** trained with  $\mathcal{L}_{Dr+}$ . To enhance the local alignment, we employ dot-regression loss  $\mathcal{L}_{DR}$ , which discards the pushing term of cross-entropy loss, and propose a feature distillation  $\mathcal{L}_{FD}$  to preserve the knowledge imbued in the global model.

align local models with the global objective through techniques like proximal terms, contrastive learning, and logit-based regularizers.

On the other direction, to improve the local alignment, recent studies have extensively focused on freezing the last linear layer (called classifier) while updating only the feature extractor. This approach is motivated by the fact that the classifier is most sensitive to data heterogeneity (Luo et al., 2021; Li et al., 2023b; Fan et al., 2024); freezing it ensures that all local models align their features to a consistent classifier across clients. For example, FedBABU (Oh et al., 2022) fixes the classifier after initialization, allowing only the feature extractor to adapt. Other methods (Dong et al., 2022; Li et al., 2023b; Fan et al., 2024; Huang et al., 2023; Xiao et al., 2024) further enhance local alignment by modifying the loss function or by using robust initialization techniques, such as the Equiangular Tight Frame (ETF) classifier.

A frozen classifier is also extensively explored in other research areas, such as class imbalance (Yang et al., 2022) and class incremental learning (Yang et al., 2023), with a consistent objective similar to aforementioned FL studies—enhancing alignment. Recently, these fields have advanced by introducing and utilizing a novel type of loss, called dot-regression loss  $\mathcal{L}_{DR}$ , which aims to achieve alignment rapidly. In summary,  $\mathcal{L}_{DR}$  originates from the decomposition analysis of cross-entropy (CE) loss, which includes *pulling* and *pushing*. As suggested in (Yang et al., 2022), the *pulling* component is a force that attracts features to the target class, whereas the *pushing* component is a force that drives features away from other non-target classes.  $\mathcal{L}_{DR}$  discards the *pushing* component, as it slows down convergence to the desired alignment (refer to Figure 1).

However, our findings indicate that while dot-regression loss enhances *local alignment* as intended, it does not lead to sufficient performance improvement of the aggregated server-side model. The main issue arises from the insufficient *global knowledge preservation* of unobserved classes during local training. Specifically, the focus on improving alignment for classes present in the local training dataset induces significant feature dynamics, which inadvertently disrupt the representation of features associated with unobserved classes. This disruption leads to forgetting and deteriorated alignment for these classes, highlighting the need for better preservation of global knowledge during local training.

We emphasize the importance of addressing both local alignment and global knowledge preservation together. By exploring the synergy between methods that target these aspects, we identified that the most effective solution is a training mechanism, termed **FedDr+**, which combines dot-regression loss with feature distillation to reduce the distance between feature vectors of local and global models.

**Contributions.** Our main contributions are summarized as follows:

- We find that dot-regression loss is not easily compatible with FL, although it can enhance the alignment of seen classes. The drawback comes from a significant loss of information on unseen classes, which is vital in the global model perspective. Therefore, we aim to preserve information of unseen classes within the FL system.
- To preserve global knowledge, including unseen class information while maintaining the advantages of  $\mathcal{L}_{DR}$ , we propose **FedDr+**, which utilizes a feature distillation when training local models. This regularizer prevents the model from focusing solely on the local alignment.
- We verify that the proposed method surpasses conventional FL algorithms under various datasets and non-iid settings.

## 2 Preliminaries

In this section, we describe the basic settings of Federated Learning (FL), including the **FedAvg** pipeline (McMahan et al., 2017) and the dot-regression loss (Yang et al., 2022), both of which are utilized in our framework. For clarity, we summarize the main notations in Table 1.

Table 1: Main Notations used throughout the paper.

<b>Indices</b>	
$c \in [C]$	Index for a class
$r \in [R]$	Index for FL round
$i \in [N]$	Index for a client
<b>Dataset</b>	
$D_{\text{train}}^i$	Training dataset for client $i$
$D_{\text{test}}^i$	Test dataset for client $i$
$(x, y) \in D_{\text{train, test}}^i ; (x, y) \sim \mathcal{D}^i$	Data on client $i$ sampled from distribution $\mathcal{D}^i$ ( $x$ : input data, $y$ : class label)
$\mathcal{O}^i$	Dataset consists of observed classes in client $i$
$\mathcal{U}^i$	Dataset consists of unobserved classes in client $i$
<b>Parameters</b>	
$\theta$	Feature extractor weight parameters
$\mathbf{V} = [v_1, \dots, v_C] \in \mathbb{R}^{C \times d}$	Classifier weight parameters (frozen during training)
$v_c, c \in [C]$	$c$ -th row vector of $\mathbf{V}$
$\Theta = (\theta, \mathbf{V})$	All model parameters
$\Theta_r^g = (\theta_r^g, \mathbf{V})$	Aggregated global model parameters at round $r$
$\Theta_r^i = (\theta_r^i, \mathbf{V})$	Trained model parameters on client $i$ at round $r$
<b>Model Forward</b>	
$p(x; \theta) \in \mathbb{R}^C$	Softmax probability of input $x$
$p_c(x; \theta), c \in [C]$	$c$ -th element of $p(x; \theta)$
$\mathcal{L}_{CE}(x; \theta) = -\log p_y(x; \theta)$	Cross-entropy loss of input $x$
$f(x; \theta) \in \mathbb{R}^d$	Feature vector of input $x$

## 2.1 Basic Setup of Conventional FedAvg Pipeline

**Basic FL setup.** Let  $[N] = \{1, \dots, N\}$  denote the indices of clients, each with a unique training dataset  $D_{\text{train}}^i = \{(x_m, y_m)\}_{m=1}^{|D_{\text{train}}^i|}$ , where  $(x_m, y_m) \sim \mathcal{D}^i$  for the  $i^{\text{th}}$  client,  $x_m$  is the input data, and  $y_m \in [C]$  is the corresponding label among  $C$  classes. Importantly, FL studies predominantly address the scenario where the data distributions are heterogeneous, *i.e.*,  $\mathcal{D}^i$  varies across clients. Knowledge distributed among clients is collected over  $R$  communication rounds. The general objective of FL is to train a model fit to the aggregated knowledge,  $\bigcup_{i \in [N]} \mathcal{D}^i$ . This objective can be seen as solving the optimization problem:

$$\min_{\Theta=(\theta, \mathbf{V})} \sum_{i \in [N]} \frac{|D_{\text{train}}^i|}{\sum_{j \in [N]} |D_{\text{train}}^j|} \mathbb{E}_{(x,y) \sim \mathcal{D}^i} [\mathcal{L}(x, y; \theta, \mathbf{V})],$$

where  $\mathcal{L}$  is the instance-wise loss function,  $\theta$  is the weight parameter for the feature extractor, and  $\mathbf{V} = [v_1, \dots, v_C] \in \mathbb{R}^{d \times C}$  is the classifier weight matrix. We use the notation  $\Theta$  to denote the entire set of model parameters.

At the beginning of each round  $r \in [R]$ , the server has access to only a subset of clients  $\mathcal{S}_r \subset [N]$  participating in the  $r^{\text{th}}$  round. At each round  $r$ , the server transmits the global model parameters  $\Theta_{r-1}^g$  to the participating clients. Each client then updates the parameters with their private data  $D_{\text{train}}^i$  and uploads  $\Theta_r^i$  to the global server. By incorporating the locally trained weights, the server then updates the global model parameters to  $\Theta_r^g$ .

**FedAvg pipeline.** Our study follows the conventional FedAvg (McMahan et al., 2017) framework to address the FL problem. FedAvg updates the global model parameters from locally trained parameters by aggregating these local models into  $\Theta_r^g = \sum_{i \in \mathcal{S}_r} w_r^i \Theta_r^i$ , where  $w_r^i = |D_{\text{train}}^i| / \sum_{j \in \mathcal{S}_r} |D_{\text{train}}^j|$  is the importance weight of the  $i^{\text{th}}$  client.

## 2.2 Dot-Regression Loss for Feature Alignment

**Dot-regression loss  $\mathcal{L}_{\text{DR}}$ .** This loss (Yang et al., 2022) facilitates a faster alignment of feature vectors (penultimate layer outputs)  $f(x; \theta) \in \mathbb{R}^d$  to the true class direction of  $v_y$ , reducing the cosine angle as follows:

$$\mathcal{L}_{\text{DR}}(x, y; \theta, \mathbf{V}) = \frac{1}{2} \left( \cos(f(x; \theta), v_y) - 1 \right)^2$$

where  $\cos(\text{vec}_1, \text{vec}_2)$  denotes the cosine of the angle between two vectors  $\angle(\text{vec}_1, \text{vec}_2)$ .

The main motivation is that the gradient of the cross-entropy (CE) loss for the feature vector can be decomposed into a *pulling* and *pushing* gradient, and recent work indicates that we can achieve better convergence by removing the pushing effect (Yang et al., 2022; Li & Zhan, 2021). The *pulling* gradient aligns  $f(x; \theta)$  with  $v_y$ , while the *pushing* gradient ensures  $f(x; \theta)$  does not align with  $v_c$  for all  $c \neq y$  (Appendix B details the exact form of pulling and pushing gradients). Since  $\mathcal{L}_{\text{DR}}$  directly attracts features to the true-class classifier, it drops the *pushing* gradient, thereby increasing the convergence speed for maximizing  $\cos(f(x; \theta), v_y)$ .

**Frozen ETF classifier.** Since  $\mathcal{L}_{\text{DR}}$  focuses on aligning feature vectors with the true-class classifier, the classifier is not required to be trained. Instead, we construct the classifier to satisfy the simplex Equiangular Tight Frame (ETF) condition, a constructive way to achieve maximum angular separation between class vectors (Yang et al., 2022; 2023). Concretely, we initialize the classifier weight  $\mathbf{V}$  as follows and freeze it throughout training:

$$\mathbf{V} \leftarrow \sqrt{\frac{C}{C-1}} \mathbf{U} \left( \mathbf{I}_C - \frac{1}{C} \mathbf{1}_C \mathbf{1}_C^\top \right),$$

where  $\mathbf{U} \in \mathbb{R}^{d \times C}$  is a randomly initialized orthogonal matrix. Note that each  $v_i$  in the classifier weight  $\mathbf{V}$  satisfies  $\cos(v_i, v_j) = -\frac{1}{C-1}$  for all  $i \neq j \in [C]$ <sup>1</sup>.

<sup>1</sup>This relation for cosines holds if the  $v_i$ 's are symmetrically distributed such that  $\bar{v} = \frac{1}{C} \sum_{i \in [C]} v_i = 0$ , and  $\cos(v_i, v_j)$  are all the same for  $i \neq j$ .




Figure 2: Comparison of (a) feature-classifier alignment gap and (b) feature-classifier alignment on the **observed** and **unobserved** classes test data for  $\theta_r^i$  trained with  $\mathcal{L}_{CE}$  and  $\mathcal{L}_{DR}$ .

### 3 When Dot-Regression Loss Meets FL

Given our focus on applying  $\mathcal{L}_{DR}$  to FL, we first examine its impact on FL models compared to the CE loss  $\mathcal{L}_{CE}$ . In summary, we find that while  $\mathcal{L}_{DR}$  improves alignment-related performance on **observed** class labels, it faces challenge with **unobserved** classes<sup>2</sup>, which are essential for the generalization objective. To address this issue, we propose **FedDr+**, which integrates  $\mathcal{L}_{DR}$  with a novel feature distillation loss. We then evaluate **FedDr+** by analyzing the effect of feature distillation and compare it with various FL algorithms and regularizers.

**Experimental configuration.** In this section, we conduct experiments on CIFAR-100 (Krizhevsky et al., 2009) with a shard non-iid setting ( $s=10$ ), where each client contains at most 10 classes. We additionally employ LDA setting ( $\alpha=0.1$ ) in Section 3.4. Refer to Section 4 for more details on the dataset configuration. The model is trained for 320 communication rounds, randomly selecting 10% of clients in each round, and the learning rate is decayed at 160<sup>th</sup> and 240<sup>th</sup> rounds. The experimental configuration for this section is detailed in subsection 4.1.

#### 3.1 Impact of $\mathcal{L}_{DR}$ on Local and Global Models

We analyze the average performance of local and global models trained with  $\mathcal{L}_{DR}$  compared to  $\mathcal{L}_{CE}$ , focusing on their ability to generalize. In Figure 2a–2b, we evaluate statistics on two datasets: the **observed** class set  $\mathcal{O}^i$ , containing classes in each client’s training data  $D_{\text{train}}^i$ , and the **unobserved** class set  $\mathcal{U}^i$ , representing unseen classes. Separately, Figure 3 reports the evaluation on all classes.

First, we examine the amount of change from the given global model to each local model in every communication round (Figure 2a). The alignment gap is denoted by  $\cos(f(x; \theta_r^i), v_y) - \cos(f(x; \theta_{r-1}^g), v_y)$ . We then evaluate the feature-classifier alignment  $\cos(f(x; \theta_r^i), v_y)$  of each local model on the test data (Figure 2b). Finally, we observe the test accuracy of the global model  $\theta_r^g$  (Figure 3).

**Alignment analysis of local models.** As shown in Figure 2a–2b,  $\mathcal{L}_{DR}$  outperforms  $\mathcal{L}_{CE}$  on **observed** classes in terms of alignment gap and alignment, while  $\mathcal{L}_{CE}$  achieves better results on **unobserved** classes for these metrics. The improvement on **observed** classes is attributed to  $\mathcal{L}_{DR}$ , which removes the pushing term present in  $\mathcal{L}_{CE}$  and concentrates its pulling effects on these classes. However, this design inherently overlooks **unobserved** classes, resulting in poorer performance on these classes compared to  $\mathcal{L}_{CE}$ .

**Accuracy result of global models.** Figure 3 shows that in the shard setting ( $s = 10$ ),  $\mathcal{L}_{CE}$  consistently outperforms  $\mathcal{L}_{DR}$ . This difference is due to the higher proportion of **unobserved** in this setting, where each client has access to at most 10 out of 100 classes.

<sup>2</sup>While we use the term “unobserved” in this context, it also applies to “rarely” existing classes.

As observed in the alignment analysis,  $\mathcal{L}_{\text{DR}}$  performs particularly poorly on the **unobserved**, with a significantly negative alignment gap and lower alignment compared to models trained with  $\mathcal{L}_{\text{CE}}$ , which contributes to the lower overall accuracy of the global model when using  $\mathcal{L}_{\text{DR}}$ .

**Importance of local alignment in observed classes.** While  $\mathcal{L}_{\text{DR}}$  struggles with local alignment on **unobserved** classes, leading to lower global accuracy compared to  $\mathcal{L}_{\text{CE}}$ , local alignment in **observed** classes remains critical. At the last learning rate decay round (240<sup>th</sup>), both methods achieve higher or comparable performance on **unobserved** classes than in the final round (320<sup>th</sup>). However, the final round shows higher global accuracy due to improved local alignment on **observed** classes. This underscores the need to preserve the advantages of  $\mathcal{L}_{\text{DR}}$  in **observed** classes while mitigating its degradation on **unobserved** classes.




Figure 3: Comparison of the global test accuracy of  $\theta_r^g$  on all classes trained using  $\mathcal{L}_{\text{CE}}$  and  $\mathcal{L}_{\text{DR}}$ .

### 3.2 FedDr+: $\mathcal{L}_{\text{DR}}$ With Feature Distillation for FL

We propose **FedDr+** to mitigate forgetting unobserved classes while retaining the strengths of dot-regression loss in aligning features of observed classes. Using  $\mathcal{L}_{\text{DR}}$  with the frozen classifier  $\mathbf{V}$ , **FedDr+** includes a regularizer that fully distills the global model’s feature vectors  $f(x; \theta^g) \in \mathbb{R}^d$  to the client features  $f(x; \theta)$ , to enhance generalization across all classes. The proposed loss function  $\mathcal{L}_{\text{Dr+}}$ , shown in (1), combines  $\mathcal{L}_{\text{DR}}$  with a regularizer  $\mathcal{L}_{\text{FD}}(x; \theta, \theta^g) = \frac{1}{d} \|f(x; \theta) - f(x; \theta^g)\|_2^2$ . Unless specified, we use a scaling parameter  $\beta = 0.9$  throughout the paper.

$$\mathcal{L}_{\text{Dr+}}(x, y; \theta, \theta^g, \mathbf{V}) = \beta \cdot \mathcal{L}_{\text{DR}}(x, y; \theta, \mathbf{V}) + (1 - \beta) \cdot \mathcal{L}_{\text{FD}}(x; \theta, \theta^g) \quad (1)$$

**Why feature distillation?** To address data heterogeneity in FL, various distillation methods have been explored, including model parameters (Oh et al., 2022; Li et al., 2020; He et al., 2020a; Li & Wang, 2019), logit-related measurement (Li & Wang, 2019; Lee et al., 2022; Itahara et al., 2021; Ye et al., 2024; Lin et al., 2020; Chen et al., 2019; Qian et al., 2022), and co-distillation (Chen et al., 2024; Cho et al., 2023). In contrast, we utilize the *feature* distillation (Heo et al., 2019) technique because the feature directly concerns alignment. On the other hand, logits lose information from features when projected onto a frozen ETF classifier (Heo et al., 2019; Li et al., 2017; 2023a; Ben-Baruch et al., 2022). By distilling features, we leverage the global, differentiated knowledge for each data input  $x$ . This approach aims to minimize drift towards observed classes, and hence, we expect it to enhance overall generalization.

### 3.3 Effect of Feature Distillation

Our findings from Section 3.1 indicate that  $\mathcal{L}_{\text{DR}}$  is unsuitable for the heterogeneous FL environment. This is primarily because there is a notable gap in how features align with the fixed classifier between  $\mathcal{O}^i$  and  $\mathcal{U}^i$ . To assess the effect of feature distillation ( $\mathcal{L}_{\text{FD}}$ ), which imposes a constraint on the feature distance  $\|f(x; \theta_r^i) - f(x; \theta_{r-1}^g)\|_2$  for  $x \in \mathcal{O}^i$ , we measure this distance for both  $\mathcal{O}^i$  and  $\mathcal{U}^i$  from the models trained with  $\mathcal{L}_{\text{DR}}$  and  $\mathcal{L}_{\text{Dr+}}$ . We additionally analyze the angle distance,  $\angle(f(x; \theta_r^i), f(x; \theta_{r-1}^g))$ , as it impacts feature-classifier alignment. These values are averaged over the selected client set  $\mathcal{S}_r$ .

**Feature distillation stabilizes the feature dynamics.** By adding  $\mathcal{L}_{\text{FD}}$ , as revealed in Figure 4a, the local model trained with  $\mathcal{L}_{\text{Dr+}}$  shows a reduction in feature distance for **observed** classes, compared to the model trained with  $\mathcal{L}_{\text{DR}}$ . This reduction happens even for **unobserved** classes. As demonstrated in Figure 4b, the overall decrease in feature distance leads to a reduction in feature angle distance for both class sets. In both local models trained with  $\mathcal{L}_{\text{DR}}$  and  $\mathcal{L}_{\text{Dr+}}$ , there is a trend where the angle distance is significantly larger for  $\mathcal{U}^i$  than for  $\mathcal{O}^i$  (Figure 4b). This large angle distance of  $\mathcal{U}^i$  leads to the degradation of the feature-classifier alignment.






Figure 4: We present (a) feature distance and (b) feature angle distance from  $\theta_{r-1}^g$  to  $\theta_r^i$  for **observed** and **unobserved** classes by training with  $\mathcal{L}_{DR}$  and  $\mathcal{L}_{Dr+}$ .




Figure 5: Comparison of feature-classifier alignment on the **observed** and **unobserved** classes test data for  $\theta_r^i$  trained with  $\mathcal{L}_{CE}$ ,  $\mathcal{L}_{DR}$  and  $\mathcal{L}_{Dr+}$ .




Figure 6: Comparison of the global test accuracy of  $\theta_r^g$  on all classes trained using  $\mathcal{L}_{CE}$ ,  $\mathcal{L}_{DR}$  and  $\mathcal{L}_{Dr+}$ .

**Feature distillation enhances local alignment and global accuracy.** As in Figure 5–6, our proposed algorithm, *i.e.*,  $\mathcal{L}_{Dr+}$ , improves feature-classifier alignment for both  $\mathcal{O}^i$  and  $\mathcal{U}^i$ , along with enhanced global accuracy compared to  $\mathcal{L}_{DR}$ . We attribute this improvement to the enhanced knowledge of the global model which is preserved by preventing the forgetting of previously trained knowledge. During the final convergence phase (240<sup>th</sup>-320<sup>th</sup>),  $\mathcal{L}_{Dr+}$  achieves **better feature-classifier alignment even on unobserved** classes compared to  $\mathcal{L}_{CE}$ . As a result, at the final round (320<sup>th</sup>),  $\mathcal{L}_{Dr+}$  demonstrates the best local alignment across all classes, surpassing both  $\mathcal{L}_{CE}$  and  $\mathcal{L}_{DR}$ , contributing to its global model’s superior performance. Even though the proposed regularizer demonstrates a reasonable regularizing effect, one question remains: “*Is it superior to other previously used regularizers?*”

### 3.4 Synergistic Effect with Different Types of FL Algorithms and Regularizers

We answer the above question by evaluating the synergy effect of various FL algorithms by maintaining their original training loss and incorporating specific regularizers, following the approach suggested in Equation 1. To address the issue of differing loss scales between the baseline FL algorithms and the regularizers, we thoroughly tune the coefficient  $\beta$  within the range  $\{0.1, 0.3, 0.5, 0.7, 0.9, 0.99, 0.999, 0.9999\}$ , and report the resulting performance in Table 2. The selected  $\beta$  values are detailed in Appendix D.

**FL algorithms and regularizers.** We evaluate several baseline FL algorithms, including dot-regression. The baseline also include FedAvg (McMahan et al., 2017)—using an unfrozen classifier, FedBABU (Oh et al., 2022)—extending FedAvg by freezing the classifier during local training, and SphereFed (Dong et al.,

Table 2: Synergy of various FL algorithms and regularizers. Baseline indicates training FL models without a regularizer. FD denotes feature distillation, which is the regularizer we use in **FedDr+**.

Algorithm	Sharding ( $s = 10$ )							LDA ( $\alpha = 0.1$ )						
	Baseline	+Prox	+KD	+NTD	+LD	+MOON	+FD	Baseline	+Prox	+KD	+NTD	+LD	+MOON	+FD
FedAvg	37.22	36.87	36.25	37.71	37.17	37.43	37.82	42.52	43.22	44.21	43.39	43.43	44.79	43.76
FedBABU	46.20	46.03	46.37	47.22	46.71	46.49	46.95	47.37	46.62	47.60	46.48	45.78	46.27	46.49
SphereFed	43.90	41.96	44.94	43.47	43.95	43.13	45.21	46.98	43.77	47.76	47.25	47.01	46.81	49.74
FedETF	32.42	31.87	32.76	32.65	32.25	34.30	32.77	46.27	45.71	46.67	46.16	45.91	45.98	46.47
FedGELA	29.17	28.69	29.11	28.84	29.36	28.80	30.33	27.11	29.03	28.45	29.62	29.41	28.09	29.75
Dot-Regression	42.52	41.95	47.45	48.32	47.52	44.72	<b>48.69</b>	42.72	46.35	49.47	50.36	49.28	50.36	<b>50.86</b>

2022)—enhancing feature-classifier alignment by using MSE loss between one-hot encoded labels and cosine similarity based logits. We also considered FedGELA and FedETF—both applying client-specific adaptive loss functions tailored to data distribution.

Alongside the FD regularizer, we evaluate a range of regularizers, including Prox (Li et al., 2020)—constraining the distance between local and global model parameters; MOON (Li et al., 2021a)—minimizing the angle distance between feature vectors of global and local models through contrastive learning; and several logit-based regularizers—KD (Hinton et al., 2015), NTD (Lee et al., 2022; Zhao et al., 2022), and LD (Kim et al., 2021)—keeping logit-related measurements of local models closely aligned with the global model. Specifically, KD applies softened softmax probability from the logit vector, NTD does the same but excludes the true class dimension, and LD distills the entire logit vector.

**FedDr+ shows best synergy.** Table 2 shows that **FedDr+** (dot-regression + FD) achieves the strongest performance among the tested combinations. FD performs exceptionally well when combined with dot-regression, SphereFed, and FedBABU. These baselines freeze the classifier while not using client-specific adaptive loss. Among them, the synergy is most effective in the order of dot-regression, SphereFed, and FedBABU—reflecting how each optimizes feature-classifier alignment well. FD outperforms other regularizers by effectively stabilizing feature dynamics for observed classes, which helps mitigate the misalignment and performance degradation for unobserved classes. MOON, on the other hand, prioritizes the cosine similarity between feature vectors from local and global models but fails to adequately control feature norm dynamics. Prox applies uniform regularization that is independent of specific data instances, leading to less refined control over feature dynamics. Logit-based regularizers lose effectiveness due to information loss when features are projected onto the classifier, as they focus on mitigating the dynamics of the less informative projected vectors rather than the richer original feature vectors.

For baselines like FedGELA and FedETF, which apply client-specific loss functions, none of the regularizers, including FD, lead to consistently lead to significant performance, particularly in the sharding setting. In non-frozen settings, FedAvg, the FD regularizer does not offer a significant performance boost. Under LDA, FedAvg combined with MOON outperforms FedAvg with FD, consistent with the claim of MOON (Li et al., 2021a).

## 4 Experiments and Results

In this section, we first present the experimental results of **FedDr+** in the context of global federated learning (GFL). We then analyze the elapsed time and conduct a sensitivity analysis for GFL, investigating the effects of varying local epochs, client sampling ratios, and different  $\beta$  values on the performance of **FedDr+**. Finally, we propose **FedDr+ FT**, which fine-tunes the **FedDr+** GFL model with  $\mathcal{L}_{Dr+}$ , and report the results comparing it with existing PFL methods.

### 4.1 Experimental Setup

**Dataset and models.** To simulate a realistic FL scenario involving 100 clients, we conduct extensive studies on three widely used datasets: CIFAR-10 (Krizhevsky et al., 2009), CIFAR-100 (Krizhevsky et al., 2009) and ImageNet-100 (Deng et al., 2009). We use VGG11 (Simonyan & Zisserman, 2014) for CIFAR-10,



Table 3: Accuracy comparison in the GFL setting. The entries are based on results obtained from three different seeds, indicating the mean and standard deviation of the accuracy of the global model, represented as  $X \pm \sigma$ . The best performance in each case is highlighted in **bold**.

NIID Partition Strategy: Sharding								
	CIFAR-100				CIFAR-10			ImageNet-100
	$s=10$	$s=20$	$s=50$	$s=100$	$s=2$	$s=5$	$s=10$	
FedAvg	36.63 $\pm$ 0.22	42.25 $\pm$ 1.42	45.57 $\pm$ 0.22	48.20 $\pm$ 1.36	72.08 $\pm$ 0.67	81.53 $\pm$ 0.35	82.38 $\pm$ 0.40	67.78 $\pm$ 0.41
FedProx	37.07 $\pm$ 0.21	42.35 $\pm$ 0.83	45.18 $\pm$ 1.07	47.78 $\pm$ 0.79	71.92 $\pm$ 0.51	81.29 $\pm$ 0.40	82.45 $\pm$ 0.35	67.81 $\pm$ 0.65
SCAFFOLD ( $\times 2$ )	46.08 $\pm$ 0.37	48.15 $\pm$ 1.21	49.31 $\pm$ 0.62	50.73 $\pm$ 0.42	75.49 $\pm$ 0.42	<b>84.14</b> $\pm$ 0.13	<b>85.11</b> $\pm$ 0.29	70.47 $\pm$ 0.46
MOON	36.95 $\pm$ 0.37	43.05 $\pm$ 0.27	43.95 $\pm$ 0.12	46.92 $\pm$ 0.08	67.55 $\pm$ 1.16	80.90 $\pm$ 0.26	82.62 $\pm$ 0.31	68.19 $\pm$ 0.32
FedNTD	34.05 $\pm$ 1.19	41.78 $\pm$ 0.31	46.42 $\pm$ 0.63	47.17 $\pm$ 0.32	72.21 $\pm$ 0.59	69.96 $\pm$ 17.10	81.99 $\pm$ 0.42	67.51 $\pm$ 0.25
FedExp	36.85 $\pm$ 0.11	42.49 $\pm$ 1.22	45.07 $\pm$ 0.92	48.09 $\pm$ 1.00	72.31 $\pm$ 0.60	81.41 $\pm$ 0.19	82.47 $\pm$ 0.16	63.34 $\pm$ 0.51
FedSOL	32.18 $\pm$ 0.18	41.54 $\pm$ 1.03	47.42 $\pm$ 0.76	47.70 $\pm$ 1.13	54.93 $\pm$ 3.52	75.73 $\pm$ 0.10	77.00 $\pm$ 0.41	66.61 $\pm$ 1.17
FedBABU	45.97 $\pm$ 0.48	45.53 $\pm$ 0.79	46.52 $\pm$ 0.51	46.02 $\pm$ 0.28	71.99 $\pm$ 0.52	81.07 $\pm$ 0.60	82.32 $\pm$ 0.06	68.82 $\pm$ 0.46
SphereFed	42.71 $\pm$ 0.65	48.63 $\pm$ 0.90	<b>52.16</b> $\pm$ 0.22	<b>53.41</b> $\pm$ 0.19	<b>76.33</b> $\pm$ 0.33	83.67 $\pm$ 0.18	84.36 $\pm$ 0.30	69.71 $\pm$ 0.39
FedETF	31.37 $\pm$ 0.72	42.22 $\pm$ 0.77	47.47 $\pm$ 0.67	49.00 $\pm$ 0.74	67.81 $\pm$ 0.94	80.78 $\pm$ 0.68	82.60 $\pm$ 0.46	70.81 $\pm$ 0.28
FedGELA	27.95 $\pm$ 0.81	38.63 $\pm$ 0.66	44.67 $\pm$ 0.51	47.95 $\pm$ 0.85	63.77 $\pm$ 2.34	79.05 $\pm$ 0.14	81.56 $\pm$ 0.09	67.08 $\pm$ 0.18
<b>FedDr+ (Ours)</b>	<b>48.21</b> $\pm$ 0.56	<b>50.77</b> $\pm$ 0.14	<b>52.15</b> $\pm$ 0.03	<b>52.41</b> $\pm$ 0.81	<b>76.57</b> $\pm$ 0.51	83.22 $\pm$ 0.34	84.14 $\pm$ 0.27	<b>71.47</b> $\pm$ 0.45

NIID Partition Strategy: LDA								
	CIFAR-100				CIFAR-10			ImageNet-100
	$\alpha=0.05$	$\alpha=0.1$	$\alpha=0.2$	$\alpha=0.3$	$\alpha=0.1$	$\alpha=0.2$	$\alpha=0.3$	
FedAvg	35.58 $\pm$ 1.35	42.10 $\pm$ 0.60	44.78 $\pm$ 0.72	45.73 $\pm$ 0.88	68.71 $\pm$ 1.82	77.75 $\pm$ 0.26	80.76 $\pm$ 0.51	65.11 $\pm$ 0.25
FedProx	37.07 $\pm$ 0.21	42.35 $\pm$ 0.83	45.18 $\pm$ 1.06	48.18 $\pm$ 0.51	69.00 $\pm$ 2.27	77.81 $\pm$ 0.24	80.55 $\pm$ 0.19	40.48 $\pm$ 1.28
SCAFFOLD ( $\times 2$ )	40.54 $\pm$ 0.48	46.14 $\pm$ 0.70	47.98 $\pm$ 0.93	48.06 $\pm$ 1.08	(Failed)	80.15 $\pm$ 0.29	<b>82.63</b> $\pm$ 0.23	66.84 $\pm$ 0.77
MOON	23.97 $\pm$ 1.15	30.86 $\pm$ 0.21	33.60 $\pm$ 0.62	35.54 $\pm$ 0.45	66.44 $\pm$ 3.28	77.36 $\pm$ 0.08	80.15 $\pm$ 0.22	41.25 $\pm$ 0.60
FedNTD	31.78 $\pm$ 3.14	40.41 $\pm$ 0.96	43.10 $\pm$ 2.03	43.04 $\pm$ 0.82	70.22 $\pm$ 0.40	77.16 $\pm$ 0.20	79.50 $\pm$ 0.56	64.87 $\pm$ 0.20
FedExp	34.39 $\pm$ 1.77	40.85 $\pm$ 1.32	44.47 $\pm$ 0.28	45.44 $\pm$ 0.14	70.14 $\pm$ 0.53	78.09 $\pm$ 0.21	80.40 $\pm$ 0.54	59.40 $\pm$ 0.36
FedSOL	34.49 $\pm$ 0.80	41.19 $\pm$ 0.30	43.55 $\pm$ 1.51	44.85 $\pm$ 0.54	59.51 $\pm$ 1.77	67.55 $\pm$ 0.41	70.96 $\pm$ 0.32	62.70 $\pm$ 0.89
FedBABU	41.97 $\pm$ 1.01	45.77 $\pm$ 0.28	44.28 $\pm$ 0.45	44.80 $\pm$ 0.63	65.15 $\pm$ 3.66	77.03 $\pm$ 0.25	79.91 $\pm$ 0.13	66.54 $\pm$ 0.30
SphereFed	39.56 $\pm$ 0.48	46.54 $\pm$ 0.58	<b>49.41</b> $\pm$ 0.78	49.22 $\pm$ 0.86	67.49 $\pm$ 3.49	80.05 $\pm$ 0.40	<b>82.62</b> $\pm$ 0.66	67.03 $\pm$ 0.30
FedETF	40.71 $\pm$ 0.90	45.63 $\pm$ 0.33	46.28 $\pm$ 1.05	46.69 $\pm$ 0.87	70.75 $\pm$ 0.36	77.86 $\pm$ 0.46	79.95 $\pm$ 0.34	68.98 $\pm$ 0.21
FedGELA	16.72 $\pm$ 1.91	27.12 $\pm$ 1.58	33.68 $\pm$ 0.19	36.17 $\pm$ 0.26	50.69 $\pm$ 7.55	66.04 $\pm$ 14.87	77.89 $\pm$ 0.97	55.57 $\pm$ 0.42
<b>FedDr+ (Ours)</b>	<b>45.12</b> $\pm$ 1.00	<b>49.48</b> $\pm$ 0.50	<b>50.67</b> $\pm$ 0.88	<b>51.15</b> $\pm$ 0.65	<b>72.07</b> $\pm$ 2.26	<b>80.90</b> $\pm$ 0.02	<b>82.42</b> $\pm$ 0.10	<b>70.20</b> $\pm$ 0.09

MobileNet (Howard et al., 2017) for CIFAR-100, and ResNet-18 (He et al., 2016) for ImageNet-100. The training data is distributed among the 100 clients using sharding and the LDA (Latent Dirichlet Allocation) partition strategies.

Following the convention, sharding distributes the data into non-overlapping shards of equal size, each shard encompassing  $\frac{|D_{\text{train}}|}{100 \times s}$  and  $\frac{|D_{\text{test}}|}{100 \times s}$  samples per class, where  $s$  denotes the number of shards per client. On the other hand, LDA involves sampling a probability vector from Dirichlet distribution,  $p_c = (p_{c,1}, p_{c,2}, \dots, p_{c,100}) \sim \text{Dir}(\alpha)$ , and allocating a proportion  $p_{c,k}$  of instances of class  $c \in [C]$  to each client  $k \in [100]$ . Smaller values of  $s$  and  $\alpha$  increase the level of data heterogeneity. For CIFAR-10 and CIFAR-100, we explore a range of  $s$  and  $\alpha$  values to assess the impact of different data heterogeneity levels. For ImageNet-100, we focus on experiments with  $s = 20$  and  $\alpha = 0.1$ .


**Implementation details.** In each round of communication, a random 10% of clients are selected to participate in the training process. The total number of communication rounds is set to 320. The initial learning rate and the number of local epochs for CIFAR-10, CIFAR-100, and ImageNet-100 are determined through grid searches, with the detailed process and results provided in Appendix C. The learning rate  $\eta$  is decayed by a factor of 0.1 at the 160th and 240th communication rounds.

## 4.2 Global Federated Learning Results

We compare **FedDr+** with a range of GFL algorithms, considering both non-freezing and freezing classifier approaches. Among non-freezing classifiers, **FedDr+** competes with FedAvg (McMahan et al., 2017), FedProx (Li et al., 2020), SCAFFOLD (Karimireddy et al., 2020), MOON (Li et al., 2021a), FedNTD (Lee et al., 2022), FedExp (Jhunjunwala et al., 2023), and FedSOL (Lee et al., 2024). **FedDr+** is also evaluated

Table 4: Elapsed time per round (in seconds) for various GFL algorithms.

	Local update without global model						Local update with global model					
	FedAvg	FedBABU	FedETF	SphereFed	FedGELA	FedExP	FedProx	SCAFFOLD	FedNTD	FedSOL	MOON	<b>FedDr+</b> (Ours)
Elapsed time	20.4	20.5	20.7	20.4	20.5	20.1	22.6	21.4	21.2	27.2	56.7	24.9

Figure 7: Performance of baselines and **FedDr+** on CIFAR-100 ( $\alpha=0.1$ ) with various analyses: (a) client sampling ratio, (b) the number of local epochs, and (c) sensitivity to  $\beta$ .

against freezing classifier algorithms such as FedBABU (Oh et al., 2022), SphereFed (Dong et al., 2022), FedETF (Li et al., 2023b), and FedGELA (Fan et al., 2024). Among the baseline algorithms, SCAFFOLD incurs a communication cost two times higher per round, denoted as ( $\times 2$ ). Our experiments encompass heterogeneous settings involving sharding and LDA non-IID environments.

Table 3 summarizes the accuracy comparison between various GFL methods proposed in recent literature and FedAvg under various conditions. While specific methods demonstrated effectiveness in particular scenarios, some of these underperformed relative to the robustness of FedAvg. For example, SCAFFOLD shown strong performance in the less heterogeneous sharding setting on CIFAR-10; however, it failed in model training under the highly heterogeneous LDA condition with  $\alpha = 0.1$ . Notably, **FedDr+** consistently outperformed all baselines in highly heterogeneous settings, achieving a 3.15% improvement in CIFAR-100 LDA with  $\alpha = 0.05$  and 3.17% in ImageNet-100 LDA.

### 4.3 Elapsed Time Results

We compare **FedDr+** with various GFL algorithms for the elapsed time per communication round on CIFAR-100 ( $s=10$ ). As shown in Table 4, incorporating a global model during the local update generally results in higher computation costs, leading to longer elapsed times compared to updates without a global model. **FedDr+** exhibits a slightly longer elapsed time than most other algorithms, yet still requires less time than FedSOL and MOON.

### 4.4 Sensitivity Analysis

We explore the impact of varying client sampling ratio and local epochs on performance, as well as the effect of different  $\beta$  values in **FedDr+**, as detailed in Figure 7. All experiments are conducted on MobileNet using the CIFAR-100 dataset with LDA ( $\alpha=0.1$ ).

**Effect of client sampling ratio and local epochs.** We evaluate the sensitivity of hyperparameters in **FedDr+** by comparing it to baselines under varying client sampling ratio and local epochs, starting from the default setting of client sampling ratio of 0.1 and local epoch of 3. Compared to FedAvg (without classifier freezing), FedBABU and SphereFed (all with classifier freezing) generally show performance improvements with increasing fraction ratios, but **FedDr+** consistently outperforms the baselines. The number of local epochs is crucial in FL; too few epochs result in underfitting, while too many cause client drift, degrading

Table 5: PFL accuracy comparison with MobileNet on CIFAR-100. For PFL, we denote the entries in the form of X, representing the mean of personalized accuracies across all clients derived from a single seed. The best performance in each case is highlighted in **bold**.

Algorithm	$s=10$	$s=20$	$s=100$	$\alpha=0.05$	$\alpha=0.1$	$\alpha=0.3$
Local only ( $\mathcal{L}_{CE}$ )	58.05	42.45	18.69	55.39	43.76	27.75
Local only ( $\mathcal{L}_{CE}+ETF$ )	58.01	41.62	18.92	55.34	43.37	27.87
Local only ( $\mathcal{L}_{DR}$ )	60.68	44.61	20.98	58.56	46.72	30.88
FedPer	70.67	57.27	24.30	62.67	53.43	35.68
Per-FedAvg	32.13	36.66	41.27	28.81	35.56	42.80
FedRep	63.14	51.69	26.31	57.53	49.60	37.00
Ditto	39.26	38.18	44.53	35.81	37.81	43.72
FedAvg-FT	69.81	56.13	47.66	63.37	56.79	50.12
FedBABU-FT	80.14	70.89	52.14	75.50	70.83	56.91
SphereFed-FT	81.90	71.56	55.83	73.21	70.00	60.03
FedETF-FT	53.75	52.94	51.69	52.96	53.97	51.67
FedGELA-FT	75.53	67.65	51.82	62.36	60.10	52.95
<b>FedDr+ FT (ours)</b>	<b>84.10</b>	<b>75.42</b>	<b>56.76</b>	<b>78.55</b>	<b>74.75</b>	<b>62.16</b>

global model performance. The default setting of local epochs 3 is optimal for all baselines, with **FedDr+** achieving the best performance. Although performance generally declines when deviating from this peak point, **FedDr+** remains the best or highly competitive.

**Weight ratio  $\beta$  analysis.** We analyze the effect of scaling parameter in **FedDr+** by varying  $\beta$  while keeping other hyperparameters constant. The performance is evaluated for  $\beta \in \{0, 0.1, 0.3, 0.5, 0.7, 0.9, 1.0\}$ . When  $\beta = 0$ , only feature distillation is applied, and when  $\beta = 1$ , only dot-regression is used.  $\beta \in \{0, 1\}$  are generally less effective, whereas  $\beta \in \{0.3, 0.5, 0.7, 0.9\}$  show consistently good performance, indicating a balanced approach is beneficial.

#### 4.5 Personalized Federated Learning Results

We introduce **FedDr+ FT**, inspired by prior work (Oh et al., 2022; Dong et al., 2022; Li et al., 2023b; Kim et al., 2023; Fan et al., 2024), which enhances personalization by leveraging local data to fine-tune the global federated learning (GFL) model. We fine-tune the **FedDr+** GFL model using  $\mathcal{L}_{DR+}$  to create **FedDr+ FT**, *i.e.*, 2-step approach. For a comprehensive analysis, we compare **FedDr+ FT** with existing personalized federated learning (PFL) methods, including 1-step approaches, *i.e.*, creating PFL models from scratch, such as FedPer (Arivazhagan et al., 2019), Per-FedAvg (Fallah et al., 2020), FedRep (Collins et al., 2021), and Ditto (Li et al., 2021b), as well as 2-step methods such as FedAVG-FT, FedBABU-FT (Oh et al., 2022), SphereFed-FT (Dong et al., 2022), FedETF-FT (Li et al., 2023b), and FedGELA-FT (Fan et al., 2024). Additionally, we compare these methods with various simple local models that have not undergone federated learning: (1) Local only ( $\mathcal{L}_{CE}$ ), trained with  $\mathcal{L}_{CE}$ , (2) Local only ( $\mathcal{L}_{CE}+ETF$ ), trained with  $\mathcal{L}_{CE}$  and initializing the classifier with an ETF classifier, and (3) Local only ( $\mathcal{L}_{DR}$ ), trained using  $\mathcal{L}_{DR}$ .

In Table 5, we first compare the performance of simple local models in PFL by examining  $\mathcal{L}_{DR}$  and  $\mathcal{L}_{CE}$ . While methods using  $\mathcal{L}_{CE}$  show no significant differences, utilizing  $\mathcal{L}_{DR}$  leads to substantial performance improvements in PFL across all settings. The “Local only ( $\mathcal{L}_{CE}$ )” and “Local only ( $\mathcal{L}_{CE}+ETF$ )” methods exhibit similar performance due to the nearly classwise orthogonal nature of randomly initialized classifiers (Oh et al., 2022; Saxe et al., 2013; Glorot & Bengio, 2010; He et al., 2015; Lezama et al., 2018). With a large number of classes ( $C=100$ ), the ETF classifier, which is also nearly classwise orthogonal, performs similarly to random initialization. When comparing **FedDr+ FT** with other 2-step methods, **FedDr+ FT** consistently demonstrates superior performance. This aligns with previous research (Nguyen et al., 2022; Chen et al., 2023) suggesting that fine-tuning from a well-initialized model yields better PFL performance. Additionally, compared with 1-step algorithms, **FedDr+ FT** continues to show superiority, outperforming all baseline methods across all settings.

## 5 Related Work

**Federated learning.** Federated Learning (FL) is a decentralized approach to deep learning where multiple clients collaboratively train a global model using their own datasets (McMahan et al., 2017; Li et al., 2020). This approach faces challenges due to data heterogeneity across clients, causing instability in the learning process (Karimireddy et al., 2020; Luo et al., 2021). To address this problem, strategies like classifier variance reduction in FedPVR (Li et al., 2022) and virtual features in CCVR (Luo et al., 2021) have been proposed. Additionally, it is essential to distinguish between Global Federated Learning (GFL) and Personalized Federated Learning (PFL), as these are crucial concepts in FL. GFL aims to improve a single global model’s performance across clients by addressing data heterogeneity through methods like client drift mitigation (Li et al., 2020; Karimireddy et al., 2020; Jhunjunwala et al., 2023), enhanced aggregation schemes (Wang et al., 2020a;b), and data sharing techniques using public or synthesized datasets (Lin et al., 2020; Luo et al., 2021). Otherwise, PFL focuses on creating personalized models for individual clients by decoupling feature extractors and classifiers for unique updates (Oh et al., 2022; Arivazhagan et al., 2019; Collins et al., 2021), modifying local loss functions (Fallah et al., 2020; Li et al., 2021b), and using prototype communication techniques (Tan et al., 2022; Xu et al., 2023).

**Frozen classifier in FL.** By focusing on alignment, previous studies have attempted to mitigate data heterogeneity by freezing the classifier (Oh et al., 2022; Dong et al., 2022; Li et al., 2023b). Nevertheless, these methods have yet to effectively improve the alignment between features and their corresponding classifier weights. Motivated by this, we integrated the dot-regression method into FL to achieve a better-aligned local model by freezing the classifier. Dot-regression, proposed to address class imbalance, focuses on aligning feature vectors to a fixed classifier, demonstrating superior alignment performance compared to previous approaches. However, optimizing the dot-regression loss to align feature vectors with a fixed classifier caused the local model to lose information on unobserved classes, thereby degrading global model performance. To address these issues, FedLoGe (Xiao et al., 2024) employing realignment techniques to ensure the well-aligned local model’s performance translated to the global model. Additionally, in FedGELA (Fan et al., 2024), the classifier is globally fixed as a simplex ETF while being locally adapted to personal distributions. Also, FedPAC Xu et al. (2023) addressed these challenges by leveraging global semantic knowledge for explicit local-global feature alignment. Besides alignment-focused methods, there have been various attempts to maintain good local model performance in the global model Jiang et al. (2023); An et al. (2024); Chen & Chao (2022).

**Knowledge distillation in FL.** Knowledge distillation (KD) has been widely studied in FL settings, such as in FedMD (Li & Wang, 2019) and FedDF (Lin et al., 2020), where a pretrained teacher model transfers knowledge to a student model. Additional distillation-based methods, such as FedFed (Yang et al., 2024) and co-distillation framework for PFL (Chen et al., 2024; Cho et al., 2023), have also been explored. In contrast to existing methods, we propose a loss function incorporating feature distillation to maintain the performance of both local and global models. To our knowledge, this is the first application of feature distillation in FL. This approach highlights the importance of distinguishing between GFL and PFL.

## 6 Conclusion

Motivated by the recent FL methods enhancing feature alignment with a fixed classifier, we first investigate the effects of applying dot-regression loss for FL. Since the dot-regression is the most direct method for feature-classifier alignment, we find it improves alignment and accuracy in local models but degrades the performance of the global model. This happens because local clients trained with dot-regression tend to forget classes that have not been observed. To address this, we propose **FedDr+**, combining dot-regression with a feature distillation method. By regularizing the deviation of local features from global features, **FedDr+** allows local models to maintain knowledge about all classes during training, thereby ultimately preserving general knowledge of the global model. Our method achieves top performance in global and personalized FL experiments, even when data is distributed unevenly across devices (non-IID settings).

## References

- Xuming An, Li Shen, Han Hu, and Yong Luo. Federated learning with manifold regularization and normalized update reaggregation. *Advances in Neural Information Processing Systems*, 36, 2024.
- Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019.
- Emanuel Ben-Baruch, Matan Karklinsky, Yossi Biton, Avi Ben-Cohen, Hussam Lawen, and Nadav Zamir. It’s all in the head: Representation knowledge distillation through classifier sharing. *arXiv preprint arXiv:2201.06945*, 2022.
- Hong-You Chen and Wei-Lun Chao. On bridging generic and personalized federated learning for image classification. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=I1hQbx10Kxn>.
- 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 *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=fWwFv--P0xP>.
- Wei-Chun Chen, Chia-Che Chang, and Che-Rung Lee. Knowledge distillation with feature maps for image classification. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14*, pp. 200–215. Springer, 2019.
- Zihan Chen, Howard Yang, Tony Quek, and Kai Fong Ernest Chong. Spectral co-distillation for personalized federated learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Yae Jee Cho, Jianyu Wang, Tarun Chirvolu, and Gauri Joshi. Communication-efficient and model-heterogeneous personalized federated learning via clustered knowledge transfer. *IEEE Journal of Selected Topics in Signal Processing*, 17(1):234–247, 2023.
- Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, pp. 2089–2099. PMLR, 2021.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- Xin Dong, Sai Qian Zhang, Ang Li, and HT Kung. Spherefed: Hyperspherical federated learning. In *European Conference on Computer Vision*, pp. 165–184. Springer, 2022.
- Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems*, 33:3557–3568, 2020.
- Ziqing Fan, Jiangchao Yao, Bo Han, Ya Zhang, Yanfeng Wang, et al. Federated learning with bilateral curation for partially class-disjoint data. *Advances in Neural Information Processing Systems*, 36, 2024.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.
- Chaoyang He, Murali Annamaram, and Salman Avestimehr. Group knowledge transfer: Federated learning of large cnns at the edge. *Advances in Neural Information Processing Systems*, 33:14068–14080, 2020a.
- Chaoyang He, Songze Li, Jinhyun So, Xiao Zeng, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Praneeth Vepakomma, Abhishek Singh, Hang Qiu, et al. Fedml: A research library and benchmark for federated machine learning. *arXiv preprint arXiv:2007.13518*, 2020b.



- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pp. 630–645. Springer, 2016.
- Byeongho Heo, Jeessoo Kim, Sangdoo Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1921–1930, 2019.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Chenxi Huang, Liang Xie, Yibo Yang, Wenxiao Wang, Binbin Lin, and Deng Cai. Neural collapse inspired federated learning with non-iid data, 2023.
- Sohei Itahara, Takayuki Nishio, Yusuke Koda, Masahiro Morikura, and Koji Yamamoto. Distillation-based semi-supervised federated learning for communication-efficient collaborative training with non-iid private data. *IEEE Transactions on Mobile Computing*, 22(1):191–205, 2021.
- Divyansh Jhunjhunwala, Shiqiang Wang, and Gauri Joshi. Fedexp: Speeding up federated averaging via extrapolation. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=IPrzNbddXV>.
- Meirui Jiang, Anjie Le, Xiaoxiao Li, and Qi Dou. Heterogeneous personalized federated learning by local-global updates mixing via convergence rate. In *The Twelfth International Conference on Learning Representations*, 2023.
- 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*, pp. 5132–5143. PMLR, 2020.
- Seongyeon Kim, Gihun Lee, Jaehoon Oh, and Se-Young Yun. Fedfn: Feature normalization for alleviating data heterogeneity problem in federated learning. *arXiv preprint arXiv:2311.13267*, 2023.
- Taehyeon Kim, Jaehoon Oh, NakYil Kim, Sangwook Cho, and Se-Young Yun. Comparing kullback-leibler divergence and mean squared error loss in knowledge distillation. *arXiv preprint arXiv:2105.08919*, 2021.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 and cifar-100 datasets. URL: <https://www.cs.toronto.edu/kriz/cifar.html>, 6(1):1, 2009.
- Gihun Lee, Minchan Jeong, Yongjin Shin, Sangmin Bae, and Se-Young Yun. Preservation of the global knowledge by not-true distillation in federated learning. *Advances in Neural Information Processing Systems*, 35:38461–38474, 2022.
- Gihun Lee, Minchan Jeong, Sangmook Kim, Jaehoon Oh, and Se-Young Yun. Fedsol: Stabilized orthogonal learning with proximal restrictions in federated learning, 2024.
- José Lezama, Qiang Qiu, Pablo Musé, and Guillermo Sapiro. Ole: Orthogonal low-rank embedding-a plug and play geometric loss for deep learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8109–8118, 2018.
- 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. *arXiv preprint arXiv:2212.02191*, 2022.



- Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581*, 2019.
- Jingzhi Li, Zidong Guo, Hui Li, Seungju Han, Ji-won Baek, Min Yang, Ran Yang, and Sungjoo Suh. Rethinking feature-based knowledge distillation for face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20156–20165, 2023a.
- Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10713–10722, 2021a.
- Quanquan Li, Shengying Jin, and Junjie Yan. Mimicking very efficient network for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6356–6364, 2017.
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.
- Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pp. 6357–6368. PMLR, 2021b.
- Xin-Chun Li and De-Chuan Zhan. Fedrs: Federated learning with restricted softmax for label distribution non-iid data. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pp. 995–1005, 2021.
- Zexi Li, Xinyi Shang, Rui He, Tao Lin, and Chao Wu. No fear of classifier biases: Neural collapse inspired federated learning with synthetic and fixed classifier. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5319–5329, 2023b.
- Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems*, 33:2351–2363, 2020.
- Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems*, 34:5972–5984, 2021.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pp. 109–165. Elsevier, 1989.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.
- 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. *arXiv preprint arXiv:2206.15387*, 2022.
- Jaehoon Oh, SangMook Kim, and Se-Young Yun. FedBABU: Toward enhanced representation for federated image classification. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=HuaYQfggn5u>.
- Biao Qian, Yang Wang, Hongzhi Yin, Richang Hong, and Meng Wang. Switchable online knowledge distillation. In *European Conference on Computer Vision*, pp. 449–466. Springer, 2022.
- Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Yue Tan, Guodong Long, Lu Liu, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 8432–8440, 2022.

- Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulou, and Yasaman Khazaeni. Federated learning with matched averaging. In *International Conference on Learning Representations*, 2020a. URL <https://openreview.net/forum?id=BkluqlSFDS>.
- Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33: 7611–7623, 2020b.
- Zikai Xiao, Zihan Chen, Liyinglan Liu, YANG FENG, Joey Tianyi Zhou, Jian Wu, Wanlu Liu, Howard Hao Yang, and Zuozhu Liu. Fedloge: Joint local and generic federated learning under long-tailed data. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=V3j5d0GQgH>.
- Jian Xu, Xinyi Tong, and Shao-Lun Huang. Personalized federated learning with feature alignment and classifier collaboration. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://openreview.net/forum?id=SXZr8aDKia>.
- Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? *Advances in Neural Information Processing Systems*, 35:37991–38002, 2022.
- Yibo Yang, Haobo Yuan, Xiangtai Li, Zhouchen Lin, Philip Torr, and Dacheng Tao. Neural collapse inspired feature-classifier alignment for few-shot class incremental learning. *arXiv preprint arXiv:2302.03004*, 2023.
- Zhiqin Yang, Yonggang Zhang, Yu Zheng, Xinmei Tian, Hao Peng, Tongliang Liu, and Bo Han. Fedfed: Feature distillation against data heterogeneity in federated learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- Rui Ye, Yaxin Du, Zhenyang Ni, Yanfeng Wang, and Siheng Chen. Fake it till make it: Federated learning with consensus-oriented generation. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=NY3wMJuaLf>.
- Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pp. 11953–11962, 2022.

- Appendix -

## FedDr+: Stabilizing Dot-regression with Global Feature Distillation for Federated Learning

The notations are organized in Appendix A. In Appendix B, we provide a detailed explanation of the pulling and pushing gradients of the CE loss. The experimental setup is described in Appendix C, which includes dataset descriptions, model specifications, the NIID partition, and the hyperparameter search process. Additional experimental results, including further analysis on the synergy effect and PFL, are presented in Appendix D.

### A Notations

Table 6: Notations used throughout the paper.

<b>Indices</b>	
$c \in [C]$	Index for a class
$r \in [R]$	Index for FL round
$i \in [N]$	Index for a client
<b>Dataset</b>	
$D_{\text{train}}^i$	Training dataset for client $i$
$D_{\text{test}}^i$	Test dataset for client $i$
$(x, y) \in D_{\text{train, test}}^i ; (x, y) \sim \mathcal{D}^i$	Data on client $i$ sampled from distribution $\mathcal{D}^i$ ( $x$ : input data, $y$ : class label)
$\mathcal{O}^i$	Dataset consists of observed classes in client $i$
$\mathcal{U}^i$	Dataset consists of unobserved classes in client $i$
<b>Parameters</b>	
$\theta$	Feature extractor weight parameters
$\mathbf{V} = [v_1, \dots, v_C] \in \mathbb{R}^{C \times d}$	Classifier weight parameters (frozen during training)
$v_c, c \in [C]$	$c$ -th row vector of $\mathbf{V}$
$\Theta = (\theta, \mathbf{V})$	All model parameters
$\Theta_r^g = (\theta_r^g, \mathbf{V})$	Aggregated global model parameters at round $r$
$\Theta_r^i = (\theta_r^i, \mathbf{V})$	Trained model parameters on client $i$ at round $r$
<b>Model Forward</b>	
$p(x; \theta) \in \mathbb{R}^C$	Softmax probability of input $x$
$p_c(x; \theta), c \in [C]$	$c$ -th element of $p(x; \theta)$
$\mathcal{L}_{\text{CE}}(x; \theta) = -\log p_y(x; \theta)$	Cross-entropy loss of input $x$
$f(x; \theta) \in \mathbb{R}^d$	Feature vector of input $x$
$z(x; \theta) = f(x; \theta) \mathbf{V}^\top \in \mathbb{R}^C$	Logit vector of input $x$
$z_c(x; \theta), c \in [C]$	$c$ -th element of $z(x; \theta)$

### B Preliminaries: Pulling and Pushing Feature Gradients in CE

In this section, we first calculate the classifier gradient for features and introduce the pulling and pushing effects of the cross-entropy objective.

#### B.1 Feature Gradient of $\mathcal{L}_{\text{CE}}$

We first provide two lemmas supporting Proposition 1, explaining the behavior of pulling and pushing feature gradients in the cross-entropy (CE) loss.

**Lemma 1.** For all  $c, c' \in [C]$ ,  $\frac{\partial p_{c'}(x; \boldsymbol{\theta})}{\partial z_c(x; \boldsymbol{\theta})} = \begin{cases} p_c(x; \boldsymbol{\theta}) \cdot (1 - p_c(x; \boldsymbol{\theta})) & \text{if } c = c' \\ -p_c(x; \boldsymbol{\theta}) \cdot p_{c'}(x; \boldsymbol{\theta}) & \text{else} \end{cases}$ .

*Proof.* Note that  $p(x; \boldsymbol{\theta}) = \left[ \frac{\exp(z_j(x; \boldsymbol{\theta}))}{\sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta}))} \right]_{j=1}^C \in \mathbb{R}^C$ . Then,

(i)  $c = c'$  case:

$$\begin{aligned} \frac{\partial p_c(x; \boldsymbol{\theta})}{\partial z_c(x; \boldsymbol{\theta})} &= \frac{\partial}{\partial z_c(x; \boldsymbol{\theta})} \left\{ \frac{\exp(z_c(x; \boldsymbol{\theta}))}{\sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta}))} \right\} \\ &= \frac{\exp(z_c(x; \boldsymbol{\theta})) \left( \sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta})) \right) - \exp(z_c(x; \boldsymbol{\theta}))^2}{\left( \sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta})) \right)^2} \\ &= p_c(x; \boldsymbol{\theta}) - p_c(x; \boldsymbol{\theta})^2 = p_c(x; \boldsymbol{\theta})(1 - p_c(x; \boldsymbol{\theta})). \end{aligned}$$

(ii)  $c \neq c'$  case:

$$\begin{aligned} \frac{\partial p_{c'}(x; \boldsymbol{\theta})}{\partial z_c(x; \boldsymbol{\theta})} &= \frac{\partial}{\partial z_c(x; \boldsymbol{\theta})} \left\{ \frac{\exp(z_{c'}(x; \boldsymbol{\theta}))}{\sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta}))} \right\} = \frac{-\exp(z_c(x; \boldsymbol{\theta})) \exp(z_{c'}(x; \boldsymbol{\theta}))}{\left( \sum_{i=1}^C \exp(z_i(x; \boldsymbol{\theta})) \right)^2} \\ &= -p_c(x; \boldsymbol{\theta}) p_{c'}(x; \boldsymbol{\theta}). \end{aligned}$$

□

**Lemma 2.**  $\nabla_{z(x; \boldsymbol{\theta})} \mathcal{L}_{CE}(x, y; \boldsymbol{\theta}) = p(x; \boldsymbol{\theta}) - \mathbf{e}_y$ , where  $\mathbf{e}_y \in \mathbb{R}^C$  is the unit vector with its  $y$ -th element as 1.

*Proof.*

$$\begin{aligned} \frac{\partial \mathcal{L}_{CE}(x, y; \boldsymbol{\theta})}{\partial z_c(x; \boldsymbol{\theta})} &= -\frac{\partial}{\partial z_c(x; \boldsymbol{\theta})} \log p_y(x; \boldsymbol{\theta}) = -\frac{1}{p_y(x; \boldsymbol{\theta})} \frac{\partial p_y(x; \boldsymbol{\theta})}{\partial z_c(x; \boldsymbol{\theta})} \\ &\stackrel{(\star)}{=} \begin{cases} p_c(x; \boldsymbol{\theta}) - 1 & \text{if } c = y \\ p_c(x; \boldsymbol{\theta}) & \text{else} \end{cases} = p_c(x; \boldsymbol{\theta}) - \mathbf{1}\{c = y\}. \end{aligned}$$

Note that  $(\star)$  holds by the Lemma 1. Therefore, the desired result is satisfied. □

**Proposition 1.** Given  $(x, y)$ , the gradient of the  $\mathcal{L}_{CE}$  with respect to  $f(x; \boldsymbol{\theta})$  is given by:

$$\nabla_{f(x; \boldsymbol{\theta})} \mathcal{L}_{CE}(x, y; \boldsymbol{\theta}) = -(1 - p_y(x; \boldsymbol{\theta}))v_y + \sum_{c \in [C] \setminus \{y\}} p_c(x; \boldsymbol{\theta})v_c. \quad (2)$$

*Proof.*

$$\begin{aligned}
& \nabla_{f(x;\theta)} \mathcal{L}_{\text{CE}}(x, y; \theta) \\
& \stackrel{(\clubsuit)}{=} [\nabla_{f(x;\theta)} z_1(x; \theta) | \cdots | \nabla_{f(x;\theta)} z_C(x; \theta)] \nabla_{z(x;\theta)} \mathcal{L}_{\text{CE}}(x, y; \theta) \\
& = \sum_{c=1}^C \frac{\partial \mathcal{L}_{\text{CE}}(x, y; \theta)}{\partial z_c(x; \theta)} \nabla_{f(x;\theta)} z_c(x; \theta) \\
& = \frac{\partial \mathcal{L}_{\text{CE}}(x, y; \theta)}{\partial z_y(x; \theta)} \nabla_{f(x;\theta)} z_y(x; \theta) + \sum_{c \in [C] \setminus \{y\}} \frac{\partial \mathcal{L}_{\text{CE}}(x, y; \theta)}{\partial z_c(x; \theta)} \nabla_{f(x;\theta)} z_c(x; \theta) \\
& = \frac{\partial \mathcal{L}_{\text{CE}}(x, y; \theta)}{\partial z_y(x; \theta)} v_y + \sum_{c \in [C] \setminus \{y\}} \frac{\partial \mathcal{L}_{\text{CE}}(x, y; \theta)}{\partial z_c(x; \theta)} v_c \\
& \stackrel{(\spadesuit)}{=} -(1 - p_y(x; \theta)) v_y + \sum_{c \in [C] \setminus \{y\}} p_c(x; \theta) v_c.
\end{aligned}$$

Employing the chain rule for  $(\clubsuit)$  and invoking Lemma 2 for  $(\spadesuit)$  confirms the result.  $\square$

## B.2 Physical Meaning of $\nabla_{f(x;\theta)} \mathcal{L}_{\text{CE}}(x, y; \theta)$

Note that  $\nabla_{f(x;\theta)} \mathcal{L}_{\text{CE}}(x, y; \theta)$  has two components:  $\mathbf{F}_{\text{Pull}} = (1 - p_y(x; \theta)) v_y$  and  $\mathbf{F}_{\text{Push}} = -\sum_{c \in [C] \setminus \{y\}} p_c(x; \theta) v_c$ .  $\mathbf{F}_{\text{Pull}}$  adjusts the feature vector in the positive direction of the actual class index's classifier vector  $v_y$ , guiding alignment towards  $v_y$ . Conversely,  $\mathbf{F}_{\text{Push}}$  adjusts the feature vector in the negative direction of the vectors in the not-true class set  $[C] \setminus \{y\}$ , inducing misalignment towards  $v_c$  for  $c \in [C] \setminus \{y\}$ .

## C Experimental Setup

### C.1 Code Implementation

Our implementations are conducted using the PyTorch framework. Specifically, the experiments presented in Table 3 and Table 4 are executed on a single NVIDIA RTX 3090 GPU, based on the code structure from the following repository: <https://github.com/Lee-Gihun/FedNTD>. The other parts of our study are carried out on a single NVIDIA A5000 GPU, utilizing the code framework from <https://github.com/jhoon-oh/FedBABU>.

### C.2 Datasets, Model, and Optimizer

To simulate a realistic FL scenario, we conduct extensive studies on three widely used datasets: CIFAR-10 (Krizhevsky et al., 2009), CIFAR-100 (Krizhevsky et al., 2009) and ImageNet-100 (Deng et al. (2009)). For each dataset, appropriate models are employed: VGG11 (Simonyan & Zisserman, 2014) for CIFAR-10, MobileNet (Howard et al., 2017) for CIFAR-100, and ResNet-18 (He et al., 2016) for ImageNet-100. A momentum optimizer is utilized for all experiments. The data preprocessing pipeline for the training phase includes **RandomResizedCrop**, **RandomHorizontalFlip**, and **Normalize** transformations for all datasets. During testing, only the **Normalize** transformation is applied for CIFAR-10 and CIFAR-100, while for ImageNet-100, **Resize**, **CenterCrop**, and **Normalize** are applied. Unless otherwise noted, the basic setting of our experiments follows the dataset statistics, FL scenario specifications, and optimizer hyperparameters summarized in Table 7.

Note: In terms of dataset information,  $C$  represents the number of classes in the dataset, with  $|D_{\text{train}}|$  and  $|D_{\text{test}}|$  indicating the total numbers of training and test data used, respectively. For the federated learning (FL) system specifics,  $R$  indicates the total number of FL rounds,  $r$  is the ratio of clients selected for each

Table 7: Summary of Dataset, Model, FL System, and Optimizer Specifications

Datasets	$C$	$ D_{\text{train}} $	$ D_{\text{test}} $	$N$	$R$	$r$	$E$	$B$	$m$	$\lambda$
CIFAR-10	10	50000	10000	100	320	0.1	10	50	0.9	1e-5
CIFAR-100	100	50000	10000	100	320	0.1	3	50	0.9	1e-5
ImageNet-100	100	130000	5000	100	320	0.1	5	50	0.9	1e-5

Table 8: Hyperparameters for VGG11 training on CIFAR-10.

	Feature un-normalized algorithms										Feature normalized algorithms		
Hyperparameters	FedAvg	FedBABU	FedProx	SCAFFOLD	MOON	FedNTD	FedExp	FedSOL	FedGELA		FedETF	SphereFed	<b>FedDr+</b> (Ours)
$\eta$	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01		0.05	0.55	0.35
Additional	None	None	$\mu=0.001$	None	$(\mu, \tau)=(1, 0.5)$	$(\beta, \tau)=(1, 3)$	$\epsilon=0.001$	$\rho=2.0$	None		$(\beta, \tau)=(1, 1)$	None	$\beta=0.9$

Table 9: Hyperparameters for MobileNet training on CIFAR-100.

	Feature un-normalized algorithms										Feature normalized algorithms		
Hyperparameters	FedAvg	FedBABU	FedProx	SCAFFOLD	MOON	FedNTD	FedExp	FedSOL	FedGELA		FedETF	SphereFed	<b>FedDr+</b> (Ours)
$\eta$	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1		0.5	6.5	5.0
Additional	None	None	$\mu=0.001$	None	$(\mu, \tau)=(1, 0.5)$	$(\beta, \tau)=(1, 3)$	$\epsilon=0.001$	$\rho=2.0$	None		$(\beta, \tau)=(1, 1)$	None	$\beta=0.9$

round, and  $E$  denotes the number of local epochs. Local model training utilizes a momentum optimizer where  $B$  is the batch size, and  $m$  and  $\lambda$  represent the momentum and weight decay parameters, respectively. The initial learning rate  $\eta$  is decayed by a factor of 0.1 at the 160th and 240th communication rounds. The initial learning rate  $\eta$  and the number of local epochs  $E$  were determined via extensive grid search for each algorithm, with details outlined in Appendix C.4.

### C.3 Non-IID Partition Strategies

To induce heterogeneity in each client’s training and test data ( $D_{\text{train}}^i, D_{\text{test}}^i$ ), we distribute the entire class-balanced datasets,  $D_{\text{train}}$  and  $D_{\text{test}}$ , among 100 clients using both sharding and Latent Dirichlet Allocation (LDA) partitioning strategies:

- **Sharding** (McMahan et al., 2017; Oh et al., 2022): We organize the  $D_{\text{train}}$  and  $D_{\text{test}}$  by label and divide them into non-overlapping shards of equal size. Each shard encompasses  $\frac{|D_{\text{train}}|}{100 \times s}$  and  $\frac{|D_{\text{test}}|}{100 \times s}$  samples of the same class, where  $s$  denotes the number of shards per client. This sharding technique is used to create  $D_{\text{train}}^i$  and  $D_{\text{test}}^i$ , which are then distributed to each client  $i$ , ensuring that each client has the same number of training and test samples. The data for each client is disjoint. As a result, each client has access to a maximum of  $s$  different classes. Decreasing the number of shards per user  $s$  increases the level of data heterogeneity among clients.
- **Latent Dirichlet Allocation (LDA)** (Luo et al., 2021; Wang et al., 2020a): We utilize the LDA technique to create  $D_{\text{train}}^i$  from  $D_{\text{train}}$ . This involves sampling a probability vector  $p_c = (p_{c,1}, p_{c,2}, \dots, p_{c,100}) \sim \text{Dir}(\alpha)$  and allocating a proportion  $p_{c,k}$  of instances of class  $c \in [C]$  to each client  $k \in [100]$ . Here,  $\text{Dir}(\alpha)$  represents the Dirichlet distribution with the concentration parameter  $\alpha$ . The parameter  $\alpha$  controls the strength of data heterogeneity, with smaller values leading to stronger heterogeneity among clients. For  $D_{\text{test}}^i$ , we randomly sample from  $D_{\text{test}}$  to match the class frequency of  $D_{\text{train}}^i$  and distribute it to each client  $i$ .

### C.4 Hyperparameter Search for $\eta$ and $E$

To optimize the initial learning rate ( $\eta$ ) and the number of local epochs ( $E$ ) for our algorithm, we conduct a grid search on the CIFAR-10, CIFAR-100, and ImageNet-100 datasets. The process and reasoning are outlined below.



Table 10: Hyperparameters for ResNet-18 training on ImageNet-100.

	Feature un-normalized algorithms										Feature normalized algorithms		
Hyperparameters	FedAvg	FedBABU	FedProx	SCAFFOLD	MOON	FedNTD	FedExP	FedSOL	FedGELA		FedETF	SphereFed	<b>FedDr+</b> (Ours)
$\eta$	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1		1.0	1.0	1.0
Additional	None	None	$\mu=0.001$	None	$(\mu, \tau)=(1, 0.5)$	$(\beta, \tau)=(1, 3)$	$\epsilon=0.001$	$\rho=2.0$	None		$(\beta, \tau)=(1, 1)$	None	$\beta=0.9$

#### C.4.1 Rationale for Varying Initial Learning Rate ( $\eta$ )

The algorithms used in our experiments differ in handling feature normalization within the loss function. Some algorithms apply feature normalization, while others do not. When features  $f(x; \theta)$  are normalized, the resulting gradient is scaled by  $\frac{1}{\|f(x; \theta)\|_2}$ . This scaling effect necessitates a grid search across various learning rates to account for the differences in learning behavior.

#### C.4.2 Rationale for Varying Local Epochs $E$

In FL, choosing the appropriate number of local epochs is crucial. Too few epochs can lead to underfitting, while too many can cause client drift. Therefore, finding the optimal number of local epochs is essential by exploring a range of values.

#### C.4.3 Grid Search Process and Results

Considering the above reasons, we perform grid search for  $\eta$  and  $E$  on CIFAR-10, CIFAR-100, and ImageNet-100 datasets. The grid search for CIFAR-10 uses a shard size of 2, while for CIFAR-100, a shard size of 10 is used. Additionally, for ImageNet-100, a shard size of 20 is used. The detailed procedures for each dataset are provided below. These optimal settings have also been confirmed to yield good performance in less heterogeneous settings.

**CIFAR-10.** We examine  $\eta$  values from  $\{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6\}$ . For  $E$ , we consider  $\{1, 3, 5, 10, 15\}$ . A default initial learning rate of 0.01 is used unless specified otherwise. The optimal learning rates vary by algorithm, and the results are summarized in Table 8. Table 8 also includes the additional hyperparameters used for each algorithm. The notation for these additional hyperparameters follows the conventions used throughout this paper (Li et al., 2020; Lee et al., 2022; Jhunjunwala et al., 2023; Li et al., 2023b; Lee et al., 2024). The optimal number of local epochs is found to be 10 for every algorithm.

**CIFAR-100.** We examine  $\eta$  values from  $\{0.1, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0, 5.5, 6.0, 6.5, 7.0\}$ . For  $E$ , we consider  $\{1, 3, 5, 10\}$ . A default initial learning rate of 0.1 is used unless specified otherwise. The optimal learning rates differ by algorithm, and the results are listed in Table 9. Table 9 also includes the additional hyperparameters used for each algorithm. The notation for these additional hyperparameters follows the conventions used throughout this paper (Li et al., 2020; Lee et al., 2022; Jhunjunwala et al., 2023; Li et al., 2023b; Lee et al., 2024). The optimal number of local epochs is found to be 3 for every algorithm.

**ImageNet-100.** We examine  $\eta$  values from  $\{0.01, 0.1, 1.0, 10.0\}$ , which are chosen to maintain a consistent logarithmic scale difference. A default initial learning rate of 0.1 is used unless specified otherwise. The optimal learning rates differ by algorithm, and the results are listed in Table 10. Table 10 also includes the additional hyperparameters used for each algorithm. The notation for these additional hyperparameters follows the conventions used throughout this paper (Li et al., 2020; Lee et al., 2022; Jhunjunwala et al., 2023; Li et al., 2023b; Lee et al., 2024). The optimal number of local epochs is fixed at 5, following the setting of (Lee et al., 2024).

## D Additional Experiment Results

### D.1 Synergy Effect Details

Table 11: Synergy of various FL algorithms and regularizers. Baseline indicates training FL models without a regularizer. FD denotes feature distillation, which is the regularizer we use in **FedDr+**.

Algorithm	Sharding ( $s = 10$ )							LDA ( $\alpha = 0.1$ )						
	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD
FedAvg	37.22	36.87	37.43	36.25	37.71	37.17	37.82	42.52	43.22	44.79	44.21	43.39	43.43	43.76
FedBABU	46.20	46.03	46.49	46.37	47.22	46.71	46.95	47.37	46.62	46.27	47.60	46.48	45.78	46.49
SphereFed	43.90	41.96	43.13	44.94	43.47	43.95	45.21	46.98	43.77	46.81	47.76	47.25	47.01	49.74
FedETF	32.42	31.87	34.30	32.76	32.65	32.25	32.77	46.27	45.71	45.98	46.67	46.16	45.91	46.47
FedGELA	29.17	28.69	28.80	29.11	28.84	29.36	30.33	27.11	29.03	28.09	28.45	29.62	29.41	29.75
Dot-Regression	42.52	41.95	44.72	47.45	48.32	47.52	<b>48.69</b>	42.72	46.35	50.36	49.47	50.36	49.28	<b>50.86</b>

Table 12: Optimal  $\beta$  value selected through grid search to achieve the best synergy of various FL algorithms and regularizers.

Algorithm	Sharding ( $s = 10$ )							LDA ( $\alpha = 0.1$ )						
	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD
FedAvg	None	0.999	0.5	0.9999	0.9999	0.999	0.9	None	0.999	0.99	0.999	0.99	0.999	0.9999
FedBABU	None	0.9999	0.9	0.999	0.99	0.999	0.999	None	0.999	0.999	0.999	0.999	0.99	0.9999
SphereFed	None	0.9999	0.9999	0.9999	0.9	0.99	0.9	None	0.9999	0.999	0.999	0.9999	0.999	0.99
FedETF	None	0.999	0.3	0.5	0.999	0.9	0.9	None	0.9999	0.9999	0.5	0.999	0.99	0.99
FedGELA	None	0.9999	0.7	0.5	0.7	0.5	0.7	None	0.99	0.9	0.5	0.5	0.5	0.3
Dot-Regression	None	0.9999	0.9	0.5	0.5	0.5	0.9	None	0.9999	0.5	0.5	0.5	0.5	0.9

Table 13: Synergy of various FL algorithms and regularizers at  $\beta = 0.9$ .

Algorithm	Sharding ( $s = 10$ )							LDA ( $\alpha = 0.1$ )						
	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD	Baseline	+Prox	+MOON	+KD	+NTD	+LD	+FD
FedAvg	37.22	30.27	36.67	35.14	35.56	34.83	37.82	42.52	36.09	42.09	41.48	41.34	43.36	43.10
FedBABU	46.20	36.71	46.49	45.50	45.09	45.81	45.31	47.37	39.04	45.92	45.58	45.56	46.46	44.77
SphereFed	43.90	1.36	1.89	41.01	43.47	41.73	45.21	46.98	1.46	2.21	45.22	46.25	43.84	48.61
FedETF	32.42	25.18	32.58	32.76	31.98	32.25	32.77	46.27	34.92	45.38	44.94	45.77	44.36	45.92
FedGELA	29.17	25.52	28.57	28.84	28.67	28.37	29.07	27.11	26.84	28.09	27.78	28.27	27.97	27.60
Dot-Regression	42.52	5.42	44.72	46.60	45.78	47.52	<b>48.69</b>	42.72	7.47	30.69	48.19	33.08	49.09	<b>50.79</b>

We evaluate the synergy effect of various FL algorithms by maintaining their original training loss while incorporating specific regularizers, as detailed in Equation 1 of the main text. To manage the differing loss scales between the baseline FL algorithms and the regularizers, we systematically tune the coefficient  $\beta$  across a range of values (0.1, 0.3, 0.5, 0.7, 0.9, 0.99, 0.999, 0.9999). The resulting performance and optimal  $\beta$  values are shown in Table 11 and Table 12. However, when we set  $\beta = 0.9$  without addressing the issue of differing loss scales, the performance results, presented in Table 13, reveal that several synergies are significantly inferior due to this oversight.

Table 14: PFL accuracy comparison with MobileNet on CIFAR-100. For PFL, we denote the entries in the form of  $X_{\pm(\gamma)}$ , representing the mean and standard deviation of personalized accuracies across all clients derived from a single seed.

Algorithm	$s=10$	$s=20$	$s=100$	$\alpha=0.05$	$\alpha=0.1$	$\alpha=0.3$
Dot-Regression	42.52	49.02	52.86	$30.31 \pm 7.95$	$37.52 \pm 5.60$	$47.08 \pm 3.69$
Dot-Regression FT ( $\mathcal{L}_{DR}$ )	$80.84 \pm (5.99)$	$74.18 \pm (5.78)$	$56.84 \pm (5.04)$	$72.02 \pm (6.80)$	$66.96 \pm (5.36)$	$60.34 \pm (3.66)$
Dot-Regression FT ( $\mathcal{L}_{DR+}$ )	$80.82 \pm (6.12)$	$73.73 \pm (5.75)$	$56.69 \pm (4.95)$	$71.85 \pm (7.03)$	$66.59 \pm (5.32)$	$59.87 \pm (3.65)$
<b>FedDr+</b> (ours)	48.69	51.00	53.23	$39.63 \pm 9.12$	$45.83 \pm 6.18$	$48.04 \pm 3.44$
<b>FedDr+ FT (<math>\mathcal{L}_{DR}</math>) (ours)</b>	<b><math>84.23 \pm (5.44)</math></b>	<b><math>75.73 \pm (4.79)</math></b>	<b><math>56.90 \pm (4.85)</math></b>	<b><math>78.65 \pm (6.17)</math></b>	<b><math>74.86 \pm (4.77)</math></b>	<b><math>62.47 \pm (3.72)</math></b>
<b>FedDr+ FT (<math>\mathcal{L}_{DR+}</math>) (ours)</b>	$84.10 \pm (5.43)$	$75.42 \pm (4.80)$	$56.76 \pm (4.91)$	$78.55 \pm (6.16)$	$74.75 \pm (4.75)$	$62.16 \pm (3.73)$

## D.2 Personalized Federated Learning Results

We introduce **FedDr+** FT and dot-regression FT, inspired by prior work (Oh et al., 2022; Dong et al., 2022; Li et al., 2023b; Kim et al., 2023). These methods enhance personalization by leveraging local data to fine-tune the GFL model. We investigate the impact of fine-tuning using  $\mathcal{L}_{DR+}$  and  $\mathcal{L}_{DR}$  loss for each GFL model to assess their effectiveness on personalized accuracy. Performance metrics without standard deviations indicate results on  $D_{\text{test}}$ , obtained from the GFL model after the initial step in the 2-step method. Our experiments involve heterogeneous settings with sharding and LDA non-IID environments, using MobileNet on CIFAR-100 datasets. We set  $s$  as 10, 20, and 100, and the LDA concentration parameter ( $\alpha$ ) as 0.05, 0.1, and 0.3. Table 14 provides detailed personalized accuracy results.

Our 2-step process involves first developing the GFL model either using dot-regression or **FedDr+**. In the second step, we fine-tune this model to create the PFL model, again using  $\mathcal{L}_{DR}$  or  $\mathcal{L}_{DR+}$ . This results in four combinations: Dot-Regression FT ( $\mathcal{L}_{DR}$ ), Dot-Regression FT ( $\mathcal{L}_{DR+}$ ), **FedDr+ FT ( $\mathcal{L}_{DR}$ )**, and **FedDr+ FT ( $\mathcal{L}_{DR+}$ )**. When the GFL model is fixed, using  $\mathcal{L}_{DR}$  for fine-tuning consistently outperforms  $\mathcal{L}_{DR+}$  across all settings, because dot-regression focuses on local alignment which advantages personalized fine-tuning. Conversely, when the fine-tuning method is fixed, employing  $\mathcal{L}_{DR+}$  for the GFL model consistently outperforms  $\mathcal{L}_{DR}$  across all settings. This aligns with previous research (Nguyen et al., 2022; Chen et al., 2023) suggesting that fine-tuning from a well-initialized model yields better PFL performance.