

Federated Learning on Heterogeneous and Long-Tailed Data via Classifier Re-Training with Federated Features

Xinyi Shang¹, Yang Lu^{1*}, Gang Huang² and Hanzi Wang¹

¹Fujian Key Laboratory of Sensing and Computing for Smart Cities, School of Informatics, Xiamen University, Xiamen, China

²Zhejiang Lab, Hangzhou, China

shangxinyi@stu.xmu.edu.cn, luyang@xmu.edu.cn, huanggang@zju.edu.cn, hanzi.wang@xmu.edu.cn

Abstract

Federated learning (FL) provides a privacy-preserving solution for distributed machine learning tasks. One challenging problem that severely damages the performance of FL models is the co-occurrence of data heterogeneity and long-tail distribution, which frequently appears in real FL applications. In this paper, we reveal an intriguing fact that the biased classifier is the primary factor leading to the poor performance of the global model. Motivated by the above finding, we propose a novel and privacy-preserving FL method for heterogeneous and long-tailed data via Classifier Re-training with Federated Features (CReFF). The classifier re-trained on federated features can produce comparable performance as the one re-trained on real data in a privacy-preserving manner without information leakage of local data or class distribution. Experiments on several benchmark datasets show that the proposed CReFF is an effective solution to obtain a promising FL model under heterogeneous and long-tailed data. Comparative results with the state-of-the-art FL methods also validate the superiority of CReFF. Our code is available at <https://github.com/shangxinyi/CReFF-FL>.

1 Introduction

The emergence of federated learning (FL) enables multiple clients to collaboratively learn a powerful global model without transmitting local private data to the server. Serving as a communication-efficient and privacy-preserving learning framework, FL has shown its potential to facilitate real-world applications in multiple domains, e.g., natural language processing [Jiang *et al.*, 2021] and fraudulent credit card detection [Zheng *et al.*, 2020]. However, during FL model training, one major practical challenge is data heterogeneity [Li *et al.*, 2020a]. In the FL environment, data generated or collected by different clients may come from various sources, which results in distribution discrepancy between clients. Moreover, real-world data often exhibits long-tail distribution with heavy class imbalance, where the number of samples in

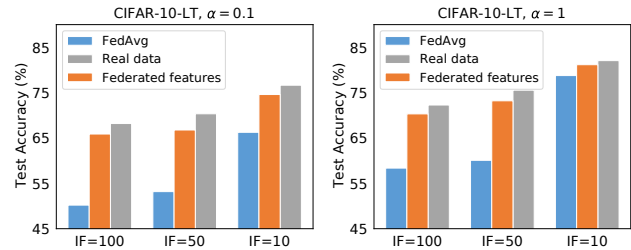


Figure 1: The performance of re-training the classifier of a FedAvg model using the same amount of data but from different sources. Real data is directly collected from clients, and federated features are synthesized by the proposed CReFF.

some classes (called head classes) severely outnumbers that in some other classes (called tail classes). Building an unbiased classification model on data with this kind of distribution is termed long-tail learning, which has been extensively studied in recent years [Zhang *et al.*, 2021]. In the FL scenario, if the training data across clients is long-tailed and heterogeneous at the same time, the joint problem becomes complicated and challenging because each client may hold different tail classes. For example, medical institutions aim to build a diagnosis model on private patient records, which are held by each institution locally. Each institution has different disease distributions, and different diseases make up long-tail distribution because some diseases are common and others are rare. Adopting FL in this case, local models will poorly perform on different tail classes due to data heterogeneity and long-tail distribution, and the aggregated global model will also be affected. One straightforward way is to simply adopt existing long-tail learning methods to the joint problem in FL. However, most of these methods require the information of local or global class distribution as prior knowledge for optimization, which may expose potential privacy issues.

One of the recent advances in long-tail learning is decoupling the training process into representation learning and classifier re-training [Kang *et al.*, 2020]. That is, with the trained feature extractor fixed, a biased classifier is re-trained using a set of balanced data, which is also called two-stage learning. In the FL framework, an FL model can also be decoupled into two-stage training, where the feature extractor is obtained by aggregating client models and the classifier is re-

*Corresponding author: Yang Lu, luyang@xmu.edu.cn

trained on the server using a set of balanced data. We conduct a quick experiment on CIFAR-10-LT (the long-tailed version of CIFAR-10 [Cao *et al.*, 2019]) to show that two-stage learning performs well on heterogeneous and long-tailed data in FL. We use imbalance factors (IF) and the Dirichlet distribution coefficient α to control the degree of long-tail and heterogeneity, respectively. Higher IF means a higher degree of imbalance degree, and lower α means a higher degree of heterogeneity. Given a global model pre-trained on heterogeneous and long-tailed data, we fix its feature extractor and re-train its classifier on the server with a small balanced dataset (100 per class) collected from clients. It can be observed in Figure 1 that the strategy is particularly useful and the classification performance is significantly improved. In particular, re-training classifier achieves the highest performance gain of 18% when IF = 100 and $\alpha = 0.1$ (the rightmost grey bar in the left subfigure of Figure 1), which is an extremely heterogeneous and long-tailed setting. However, this prerequisite of using a shared balanced dataset on the server is infeasible for most FL practical scenarios due to privacy concerns.

Motivated by the idea of two-stage learning and the privacy concerns of FL, we propose a novel and privacy-preserving FL method called Classifier Re-training with Federated Features (CReFF) to deal with the joint problem of data heterogeneity and long-tail distribution. We learn a set of balanced features called *federated features* because only the classifier needs to be re-trained on the server. It is based on an intuitive idea that the classifier re-trained on federated features should produce comparable performance as the one re-trained on the real data, which can be achieved by making two classifiers similar. Therefore, the classifier re-training optimization on federated features should follow the similar path as that of the real data. Specifically, we optimize the federated features to make their gradients close to the gradients of real data. As shown in Figure 1, the classifier re-trained on federated features (orange bars) achieves comparable performance as the one re-trained on real data (grey bars). Extensive experiments show that CReFF significantly outperforms the state-of-the-art federated learning methods on image classification tasks with heterogeneous and long-tailed data distribution. The contributions of this paper can be summarized as follows:

- We study the joint problem of FL with heterogeneous and long-tailed data distribution, where the local and global class distribution is unknown to the server.
- We reveal an intriguing fact that the biased classifier is the primary factor leading to the poor performance of an FL global model on heterogeneous and long-tailed data.
- We propose CReFF, a novel FL algorithm to deal with heterogeneous and long-tailed data by re-training the classifier with learnable federated features on the server. CReFF has no privacy concerns because no real data or information of class distribution is required.

2 Related Work

2.1 Federated Learning with Heterogeneous Data

A variety of solutions have been proposed to tackle data heterogeneity in FL, which are mainly from two perspectives:

One focuses on optimization strategies to make the diversity between client models and global model limited [Li *et al.*, 2020b; Huang *et al.*, 2021]; the other adopt mechanisms on the server to alleviate the negative influence of data heterogeneity [Lin *et al.*, 2020; Chen and Chao, 2021]. For example, CCVR [Luo *et al.*, 2021] deals with data heterogeneity via classifier re-training using virtual features sampled from an approximated Gaussian mixture model. Although the abovementioned methods solve the data heterogeneity to some extent, they poorly perform on the tail classes due to the lack of consideration of the global long-tail distribution.

Some FL methods are designed for imbalanced data (not specifically long-tailed). One strategy is to adopt client selection to match clients with complementary class distribution [Duan *et al.*, 2020; Yang *et al.*, 2020]. In this case, all clients are required to upload their local class distributions to the server, which violates the principle of privacy protection of FL. Ratio loss [Wang *et al.*, 2021a] utilizes balanced auxiliary data on the server to estimate the global class distribution for better local optimization. However, as we mentioned above, the auxiliary data is not available on the server in real applications. Compared with these related methods, the proposed CReFF has no privacy concerns because no real data or information of class distribution is required.

2.2 Long-tail Learning

Recently, long-tail learning has drawn much interest in deep learning [Zhang *et al.*, 2021]. Some methods follow the ideas of imbalance learning to augment the feature space for rare classes [Kim *et al.*, 2020; Zang *et al.*, 2021] or re-weight different classes according to their frequencies [Cao *et al.*, 2019]. Some recent methods decouple the training phase into representation learning and classifier re-training [Kang *et al.*, 2020], which generate more generalizable representations and achieve strong performance after re-balancing the classifier. Based on decoupling, ensemble-based methods [Xiang *et al.*, 2020; Wang *et al.*, 2021b] adopt a multi-expert framework to learn diverse classifiers in parallel, which reduces the model bias towards the tail classes. However, most of them require the global class distribution, as we mentioned above. During FL model training, it is infeasible to gather the information of class distribution of each client to obtain the global class distribution, which makes the vast majority of long-tail learning methods not applicable to FL scenarios.

3 Proposed Method

3.1 Preliminaries

Settings and Notations. We discuss a typical FL setting with K clients holding potentially heterogeneous data partition $\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K$, respectively. The goal is to learn a global model over the whole training data $\mathcal{D} \triangleq \bigcup_k \mathcal{D}^k$ with the help of a central server without data transmitting. In this paper, we consider the setting when \mathcal{D} is drawn from a long-tail distribution $\mathcal{X} = \{\mathbf{x}_i, y_i\}_{i=1}^N, y_i \in \{1, \dots, C\}$. Let n_c^k be the number of samples of class c on client k , and $n_c = \sum_{k=1}^K n_c^k$. Without loss of generality, a common assumption in long-tail learning is that the classes are sorted by cardinality in non-increasing order, i.e., if $c_1 < c_2$, we have

$n_{c_1} \geq n_{c_2}$ and $n_1 \gg n_C$. For an FL model, we typically consider a neural network $\phi_{\mathbf{w}}$ with parameters $\mathbf{w} = \{\mathbf{u}, \mathbf{v}\}$. It has two main components: 1) a feature extractor $f_{\mathbf{u}}$ with parameters \mathbf{u} , mapping each input sample \mathbf{x} to a d -dim feature vector; 2) a classifier $h_{\mathbf{v}}$ with parameters \mathbf{v} , typically being a fully-connected layer which outputs logits to denote class confidence scores. The parameters of the k th client local model are denoted as \mathbf{w}_k .

Basic Algorithm of FL. FedAvg [McMahan *et al.*, 2017] is the fundamental algorithm for FL. In round t , the server first sends a global model \mathbf{w}^t to clients. The clients then update the received model on their local data \mathcal{D}^k , $k = 1, \dots, K$:

$$\mathbf{w}_k^{t+1} \leftarrow \mathbf{w}_k^t - \eta \nabla_{\mathbf{w}} \ell(\mathbf{w}^t; \mathcal{D}^k). \quad (1)$$

After local updating, some clients in the set \mathcal{A}^t are selected to upload their updated models to the server. Finally, the server performs weighted average to update the global model for round $t + 1$:

$$\mathbf{w}^{t+1} = \sum_{k \in \mathcal{A}^t} \frac{|\mathcal{D}^k|}{\sum_{k \in \mathcal{A}^t} |\mathcal{D}^k|} \mathbf{w}_k^{t+1}. \quad (2)$$

3.2 Motivation

Classifier re-training has been shown effective for heterogeneous data [Luo *et al.*, 2021] and long-tailed data [Kang *et al.*, 2020], separately. Now, we empirically show that classifier re-training also works for the joint problem. In other words, the biased classifier is the primary factor leading to the poor performance of the global model trained on heterogeneous and long-tailed data.

First, we divide a ResNet-8 network into five components, including four blocks¹ and one classifier. Then, we use a small set of balanced data (100 per class) collected from clients to re-train each component of a FedAvg model while keeping other components fixed. The FedAvg model is pre-trained on CIFAR-10-LT with $\text{IF} = 100$ and $\alpha = 0.5$. As shown in Figure 2, we can observe that re-training any single component can achieve certain improvement (blue line). Especially, re-training the classifier achieves the highest performance gain (around 15.35%). Therefore, it can be shown that the biased classifier is the primary factor leading to the poor performance of the global model because other components are less affected.

However, requiring external data on the server is impractical in real FL applications, although the amount of the data is small. Transmitting data from clients violates the key privacy-preserving principle of FL, and the data collected by the server itself may not follow the same distribution as the training data in clients.

3.3 Framework of the Proposed CReFF

Based on the above assumption and corresponding observation, we propose CReFF to deal with the joint problem of data heterogeneity and long-tail distribution. It is designed as a simple and effective approach based on FedAvg by only

¹We regard the first convolutional layer with BN layer as the first block.




Figure 2: The performance of re-training each component of a FedAvg model pre-trained on CIFAR-10-LT with $\text{IF} = 100$ and $\alpha = 0.5$.

re-training a new classifier using a small set of learnable features on the server, which is called *federated features*. Our goal is to re-train the classifier on federated features such that its performance is comparable with the classifier re-trained on real features, which can be achieved by making two classifiers similar. Therefore, the classifier re-training optimization on federated features should follow the similar path as that of the real data. Specifically, we optimize the federated features to make their gradients close to the gradients of real data such that the classifiers trained on real and federated features converge to a similar solution in the parameter space. In this way, the classifier re-trained on the federated features can be approximated by the one re-trained on the real data in a privacy-preserving manner. As shown in Figure 2, CReFF produces comparable performance as on real data.

CReFF consists of two core components: local training on clients and federated features optimization on the server. In round t , each client receives two models from the server: 1) a global model \mathbf{w}^t ; 2) a re-trained model $\hat{\mathbf{w}}^t$. The former is used to update local models for global model aggregation, and the latter is used to calculate the real feature gradients for federated features optimization. After local training, some clients are selected to upload their local models and real feature gradients to the server. Finally, the server updates the global model and the re-trained model for round $t + 1$. The architecture of CReFF is illustrated in Figure 3.

Local Training. In round t , the server sends two models to clients. The global model \mathbf{w}^t is composed of a feature extractor \mathbf{u}^t and a classifier \mathbf{v}^t , and the re-trained model $\hat{\mathbf{w}}^t$ has a re-trained classifier $\hat{\mathbf{v}}^t$ with the same feature extractor \mathbf{u}^t . For each received model, the local training consists of a corresponding part. The first part is a typical FedAvg local model update as shown in Equation (1). The second part is calculating the real feature gradients for federated features optimization on the server. Client k produces d -dim real features $\mathcal{Z}_c^k = \{\mathbf{z}_{c,i}^k\}_{i=1}^{n_c^k}$ for class c through the feature extractor \mathbf{u}^t . Then, real feature gradients $\mathbf{g}_c^k \in \mathbb{R}^{C \times d}$ of class c can be computed using $\hat{\mathbf{v}}^t$:

$$\mathbf{g}_c^k = \frac{1}{n_c^k} \sum_{i=1}^{n_c^k} \nabla_{\hat{\mathbf{v}}^t} \ell(h_{\hat{\mathbf{v}}^t}(\mathbf{z}_{c,i}^k), y_i). \quad (3)$$

Finally, client k uploads two parts to the server: 1) the local model \mathbf{w}_k^{t+1} for global model aggregation; 2) the real feature gradients $\{\mathbf{g}_c^k | c \in \mathcal{C}^k\}$ for federated features optimization. Herein, \mathcal{C}^k denotes the set of classes on client k because

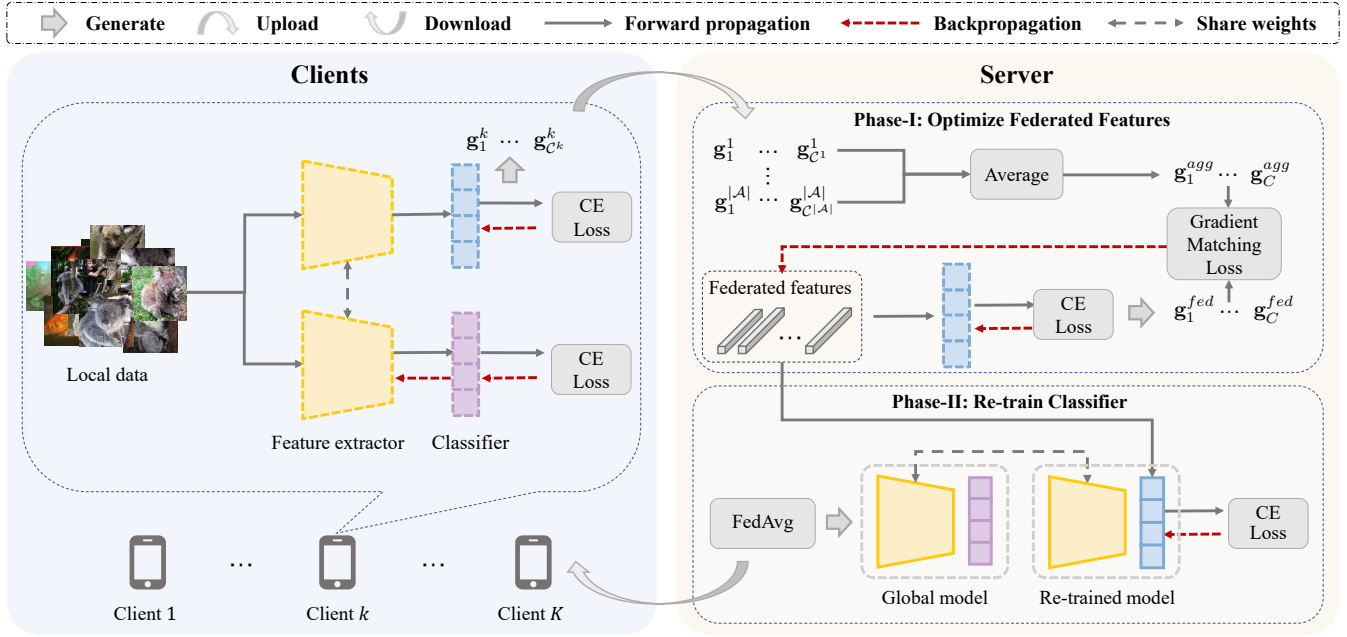


Figure 3: The framework of CReFF. In each round, clients send updated local models and real feature gradients to the server, and the server sends the aggregated global model and the re-trained model to clients.

the real feature gradients on client k may only contain partial classes due to data heterogeneity. It is worth noting that clients only upload the average gradients of each class to the server, which erases the private information of local data to some context because the averaging operation is irreversible.

Federated Features Optimization. Once the server receives two parts from the selected clients in \mathcal{A}^t , there are also two corresponding parts to aggregate the client models and optimize federated features for classifier re-training, respectively. The first part is to aggregate the local models by Equation (2) to obtain the aggregated global model \mathbf{w}^{t+1} . The second part aims to learn a set of balanced d -dim federated features $\mathcal{S}_c^t = \{\mathbf{s}_{c,i}^t\}_{i=1}^m$ for each class c in round t , where m is the number of federated features of each class. First, the server aggregates real feature gradients of each class c by averaging over all selected clients in \mathcal{A}^t :

$$\mathbf{g}_c^{agg} = \frac{1}{|\mathcal{A}_c^t|} \sum_{k=1}^{|\mathcal{A}_c^t|} \mathbf{g}_c^k, \quad (4)$$

where \mathcal{A}_c^t is the subset of clients that hold class c . Then, the server computes the federated feature gradients \mathbf{g}_c^{fed} over federated features \mathcal{S}^t of class c by using the same re-trained classifier $\hat{\mathbf{v}}^t$ as calculating \mathbf{g}_c^{agg} :

$$\mathbf{g}_c^{fed} = \frac{1}{m} \sum_{i=1}^m \nabla_{\hat{\mathbf{v}}} \ell(h_{\hat{\mathbf{v}}^t}(\mathbf{s}_{c,i}^t), y_i). \quad (5)$$

Both \mathbf{g}_c^{fed} and \mathbf{g}_c^{agg} are gradients of the same re-trained classifier $\hat{\mathbf{v}}$, whose dimensions are $C \times d$. We use the gradient matching loss [Zhao *et al.*, 2021] to measure the difference between them by averaging the cosine dissimilarity of each

part:

$$D(\mathbf{g}_c^{fed}, \mathbf{g}_c^{agg}) = \frac{1}{C} \sum_{j=1}^C \left(1 - \frac{\mathbf{g}_c^{fed}[j] \cdot \mathbf{g}_c^{agg}[j]}{\|\mathbf{g}_c^{fed}[j]\| \|\mathbf{g}_c^{agg}[j]\|} \right), \quad (6)$$

where $\mathbf{g}[j]$ denotes the j th row of the gradient. After optimizing the federated features \mathcal{S}^t on this loss, we obtain an updated \mathcal{S}^{t+1} for classifier re-training.

Compared with sampling virtual features from the approximated Gaussian mixture for classifier re-training [Luo *et al.*, 2021], CReFF learns federated features by optimizing the gradient matching loss. The advantages are twofold. On the one hand, CReFF can produce nearly the same gradients as the real features to update the classifier, which can not be achieved by CCVR. On the other hand, CCVR requires each client to upload its local class distribution to the server for computing global mean and covariance. In contrast, CReFF does not require the information of local class distribution in a more privacy-preserving manner.

The last step of CReFF is to re-train the classifier of the updated global model \mathbf{w}^{t+1} to obtain the re-trained model $\hat{\mathbf{w}}^{t+1}$. Initialized as $\tilde{\mathbf{v}}$, a new classifier $\hat{\mathbf{v}}^{t+1}$ is re-trained on the optimized balanced federated features \mathcal{S}^{t+1} :

$$\hat{\mathbf{v}}^{t+1} \leftarrow \tilde{\mathbf{v}} - \eta \nabla_{\hat{\mathbf{v}}} \ell(\tilde{\mathbf{v}}; (\mathbf{s}_i^t, y_i)). \quad (7)$$

Finally, the updated global model $\mathbf{w}^{t+1} = \{\mathbf{u}^{t+1}, \mathbf{v}^{t+1}\}$ and the re-trained model $\hat{\mathbf{w}}^{t+1} = \{\mathbf{u}^{t+1}, \hat{\mathbf{v}}^{t+1}\}$ are sent to clients for round $t+1$. The whole training process of CReFF is shown in Algorithm 1.

4 Experimental Results

In this section, we first show the ability of CReFF to deal with data heterogeneity and long-tail distribution on several

Algorithm 1: Training process of CReFF

Input: Initialized global model \mathbf{w}^0 , number of steps I for federated features optimization, number of steps J for classifier re-training, number of federated features per class m , number of training rounds T

Output: Re-trained model $\widehat{\mathbf{w}}^T$ on round T

```

1 for  $t = 1$  to  $T$  do
2   Randomly select a set of active clients  $\mathcal{A}^t$ ;
   // Clients execute:
3   for  $k \in \mathcal{A}^t$  do
4     Update local model  $\mathbf{w}_k^{t+1}$  by Equation (1);
5     Compute real feature gradients  $\{\mathbf{g}_c^k | c \in \mathcal{C}^k\}$ 
     by Equation (3);
6     Send  $\mathbf{w}_k^{t+1}$  and  $\{\mathbf{g}_c^k | c \in \mathcal{C}^k\}$  to the server;
7   end
   // Server executes:
8   Aggregate local models to  $\mathbf{w}^{t+1}$  by Equation (2);
9   Aggregate real feature gradients to  $\mathbf{g}_c^{agg}$  by
   Equation (4);
10  Compute federated feature gradients  $\mathbf{g}_c^{fed}$  by
   Equation (5);
11  Optimize federated features to  $\mathcal{S}^{t+1}$  with the loss
   in Equation (6) for  $I$  epochs;
12  Re-train the classifier of  $\mathbf{w}^{t+1}$  to  $\widehat{\mathbf{w}}^{t+1}$  on  $\mathcal{S}^{t+1}$  by
   Equation (7) for  $J$  epochs;
13  Send  $\mathbf{w}^{t+1}$  and  $\widehat{\mathbf{w}}^{t+1}$  to clients;
14 end
    
```

benchmark datasets, compared with the state-of-the-art FL methods. Then, we specifically study the reason of federated features being effective for classifier re-training on the server.

4.1 Experimental Setup

We conduct experiments on the following datasets:

- CIFAR-10/100-LT [Krizhevsky and Hinton, 2009]. We follow [Cao *et al.*, 2019] to shape the original balanced CIFAR-10/100 into long-tail distribution with IF = 100, 50 and 10, respectively. Like previous studies [Lin *et al.*, 2020], we use Dirichlet distribution to generate the heterogeneous data partition among clients. We set the value of α at 0.5 on CIFAR-10/100-LT.

- ImageNet-LT [Russakovsky *et al.*, 2015]. It contains 115.8K images from 1,000 categories, with the largest and smallest categories containing 1,280 and 5 images, respectively. We set the value of α at 0.1 on ImageNet-LT.

We use ResNet-8 for CIFAR-10/100-LT, and ResNet-50 for ImageNet-LT as the base model. We implement all compared FL methods with the same model for a fair comparison. All experiments are run by PyTorch on two NVIDIA GeForce RTX 3080 GPUs. By default, we use standard cross-entropy loss and run 200 communication rounds. We set the number of total clients at 20 and an active client ratio 40% in each round. For local training, the batch size is set at 32. For server-side training, we initialize the federated features as random noise and set the number of them per class at 100,

the optimization steps I on federated features at 100, the classifier re-training steps J at 300. We use SGD with a learning rate 0.1 as the optimizer for all optimization process.

4.2 Comparison with the State-of-the-art Methods

We compare CReFF with heterogeneity-oriented FL methods, including FedAvg [McMahan *et al.*, 2017], FedAvgM [Hsu *et al.*, 2019], FedProx [Li *et al.*, 2020b], FedDF [Lin *et al.*, 2020], FedBE [Chen and Chao, 2021], CCVR [Luo *et al.*, 2021] and FedNova [Wang *et al.*, 2020]. Moreover, we also compare with imbalance-oriented FL methods, including Fed-Focal Loss [Sarkar *et al.*, 2020], Ratio Loss [Wang *et al.*, 2021a] and FedAvg with τ -norm [Kang *et al.*, 2020]. Note that Ratio Loss depends on a balanced auxiliary data, so we fetch the data from clients, although it is infeasible.

Results on CIFAR-10/100-LT. The results are summarized in Table 1. CReFF achieves the highest test accuracy on both datasets with different IFs. Compared with the baseline FedAvg, the performance gain of CReFF is the highest when IF = 100 (around 14.4% improvement for CIFAR-10-LT and 4.3% improvement for CIFAR-100-LT). It shows the generalization ability of CReFF when the global class distribution is highly long-tailed. For heterogeneity-oriented methods, most of them (e.g., FedAvgM and FedDF) perform similarly to FedAvg because they only deal with data heterogeneity without taking global class imbalance into account. For imbalance-oriented methods, some of them, e.g., Ratio Loss, perform well in some cases compared with FedAvg. However, there is still a performance gap compared with CReFF because they aim at dealing with the global imbalance problem but generally ignore the data heterogeneity problem across clients.

Results on ImageNet-LT. We further evaluate CReFF on ImageNet-LT as reported in Table 2. To better examine the performance of classes with different number of training samples, we report the accuracy on three sets of classes: Many-shot (more than 100 samples), medium-shot (20-100 samples), and few-shot (less than 20 samples). Compared with other FL methods, CReFF achieves the best results on all cases, which shows that the improved overall accuracy by CReFF does not sacrifice the accuracy of many-shot classes.

4.3 Model Validation

To further validate the effectiveness of CReFF, we study two key research questions.

Why does a classifier re-trained on federated features perform similarly to the one re-trained on real features? We show the dissimilarity between the federated and the real feature gradients by the gradient matching loss of each class. We group the classes into three sets for CIFAR-10-LT: Many (more than 1500 samples), medium (200-1500 samples) and few (less than 200 samples), and then average the dissimilarities in each group shown by the solid lines in Figure 4. It can be observed that the dissimilarities decrease to around 0 in a few rounds, which means that optimizing on the gradient matching loss successfully makes the federated feature gradients close to the real feature gradients. Thus, the optimization of classifier re-training on federated features follows a similar path as that of the real features. Moreover, we also show

Family	Method	CIFAR-10-LT			CIFAR-100-LT		
		IF = 100	IF = 50	IF = 10	IF = 100	IF = 50	IF = 10
Heterogeneity-oriented FL methods	FedAvg	56.17	59.36	77.45	30.34	36.35	45.87
	FedAvgM	52.03	57.11	70.81	30.80	35.33	44.66
	FedProx	56.92	60.89	76.53	31.67	36.30	46.10
	FedDF	55.15	58.74	76.51	31.43	36.22	46.19
	FedBE	55.79	59.55	77.78	31.97	36.39	46.25
	CCVR	69.53	71.89	78.48	33.43	36.98	46.88
	FedNova	57.79	63.91	77.79	32.64	36.62	46.75
Imbalance-oriented FL methods	Fed-Focal Loss	53.83	57.42	73.74	30.67	35.25	45.52
	Ratio Loss	59.75	64.77	78.14	32.95	36.88	46.79
	FedAvg+ τ -norm	49.95	51.41	72.08	26.22	33.71	43.65
Proposed method	CRoFF	70.55	73.08	80.71	34.67	37.64	47.08

Table 1: Top-1 test accuracy (%) achieved by compared FL methods and CRoFF on CIFAR-10/100-LT with different IFs.

Method	ImageNet-LT			
	All	Many	Medium	Few
FedAvg	23.85	34.92	19.18	7.10
FedAvgM	22.57	33.93	18.55	6.73
FedProx	22.99	34.25	17.06	6.37
FedDF	21.63	31.78	15.52	4.48
CCVR	25.49	36.72	20.24	9.26
Fed-Focal Loss	21.60	31.74	15.77	5.52
Ratio Loss	24.31	36.33	18.14	7.41
FedAvg+ τ -norm	21.58	31.66	15.76	4.33
CRoFF	26.31	37.44	21.87	10.29

Table 2: Top-1 test accuracy (%) achieved by compared FL methods and CRoFF on ImageNet-LT.

the dissimilarity between the federated and the real features by calculating their cosine dissimilarity in Figure 4 (dotted lines). As training goes on, the dissimilarities between the federated and real features are also decreasing but they are far from 0. It means that the federated features do not have to be similar to the real features although their corresponding gradients are very close. Therefore, the privacy of the real features can be further guaranteed.

How many federated features per class is enough to re-train a good classifier? One important hyperparameter in CRoFF is the number of federated features per class. We evaluate its influence by tuning it in the set $\{0, 1, 5, 10, 50, 100, 150, 200\}$ on CIFAR-10-LT with different IFs. It can be observed from Figure 5 that more federated features produce higher accuracy in general. However, as the number of federated features increases, the improvement gradually becomes stable. This result is anticipated because the information provided by real feature gradients is limited. Meanwhile, a surprising result is that even if only one federated feature is learned for each class, CRoFF can still improve 4.21%, 3.74%, and 0.92% for CIFAR-10-LT with IF = 100, 50 and 10, respectively.




 Figure 4: The lines of two dissimilarities on CIFAR-10-LT with IF = 100 and $\alpha = 0.5$.


Figure 5: The performance of CRoFF with the different number of federated features on CIFAR-10-LT with different IFs.

5 Conclusion

In this paper, we first show that classifier re-training is a straightforward and effective way to deal with the joint problem of data heterogeneity and long-tail distribution in the federated learning framework. Based on this, CRoFF is proposed to learn a small set of balanced features called *federated features* on the server for classifier re-training in a privacy-preserving manner. The classifier re-trained on the federated features can produce comparable performance as the one re-trained on the real data. Experiments have shown that CRoFF outperforms the state-of-the-art FL methods in the setting of heterogeneity and long-tail distribution, and the effectiveness CRoFF is validated empirically.

Acknowledgments

We thank Dr. Fei Chen for his helpful discussion and reviewers of IJCAI-ECAI 2022 for their constructive and helpful feedback. This work was supported in part by the National Natural Science Foundation of China (No. 62002302, U21A20514, 61872307 and 52007173), the Open Research Projects of Zhejiang Lab (No. 2021KB0AB03), the Natural Science Foundation of Fujian Province (No. 2020J01005), and Zhejiang Provincial Natural Science Foundation of China (No. LQ20E070002).

References

- [Cao *et al.*, 2019] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arachiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. In *NeurIPS*, pages 1567–1578, 2019.
- [Chen and Chao, 2021] Hong-You Chen and Wei-Lun Chao. Fedbe: Making bayesian model ensemble applicable to federated learning. In *ICLR*, 2021.
- [Duan *et al.*, 2020] Moming Duan, Duo Liu, Xianzhang Chen, Renping Liu, Yujuan Tan, and Liang Liang. Self-balancing federated learning with global imbalanced data in mobile systems. *IEEE Transactions on Parallel and Distributed Systems*, 32(1):59–71, 2020.
- [Hsu *et al.*, 2019] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019.
- [Huang *et al.*, 2021] Hua Huang, Fanhua Shang, Yuanyuan Liu, and Hongying Liu. Behavior mimics distribution: Combining individual and group behaviors for federated learning. In *IJCAI*, pages 2556–2552, 2021.
- [Jiang *et al.*, 2021] Ziyue Jiang, Yi Ren, Ming Lei, and Zhou Zhao. Fedspeech: Federated text-to-speech with continual learning. In *IJCAI*, pages 3829–3835, 2021.
- [Kang *et al.*, 2020] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. In *ICLR*, 2020.
- [Kim *et al.*, 2020] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. In *CVPR*, pages 13896–13905, 2020.
- [Krizhevsky and Hinton, 2009] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. In *Technical report*, pages 32–33. University of Toronto, 2009.
- [Li *et al.*, 2020a] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- [Li *et al.*, 2020b] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *MLSys*, pages 429–450, 2020.
- [Lin *et al.*, 2020] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In *NeurIPS*, pages 2351–2363, 2020.
- [Luo *et al.*, 2021] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. In *NeurIPS*, 2021.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pages 1273–1282, 2017.
- [Russakovsky *et al.*, 2015] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [Sarkar *et al.*, 2020] Dipankar Sarkar, Ankur Narang, and Sumit Rai. Fed-focal loss for imbalanced data classification in federated learning. *arXiv preprint arXiv:2011.06283*, 2020.
- [Wang *et al.*, 2020] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. In *NeurIPS*, pages 7611–7623, 2020.
- [Wang *et al.*, 2021a] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. Addressing class imbalance in federated learning. In *AAAI*, pages 10165–10173, 2021.
- [Wang *et al.*, 2021b] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X. Yu. Long-tailed recognition by routing diverse distribution-aware experts. In *ICLR*, 2021.
- [Xiang *et al.*, 2020] Liuyu Xiang, Guiguang Ding, and Jun-gong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *ECCV*, pages 247–263, 2020.
- [Yang *et al.*, 2020] Miao Yang, Akitanoshou Wong, Hongbin Zhu, Haifeng Wang, and Hua Qian. Federated learning with class imbalance reduction. *arXiv preprint arXiv:2011.11266*, 2020.
- [Zang *et al.*, 2021] Yuhang Zang, Chen Huang, and Chen Change Loy. Fasa: Feature augmentation and sampling adaptation for long-tailed instance segmentation. In *ICCV*, 2021.
- [Zhang *et al.*, 2021] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *arXiv preprint arXiv:2110.04596*, 2021.
- [Zhao *et al.*, 2021] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In *ICLR*, 2021.
- [Zheng *et al.*, 2020] Wenbo Zheng, Lan Yan, Chao Gou, and Fei-Yue Wang. Federated meta-learning for fraudulent credit card detection. In *IJCAI*, pages 4654–4660, 2020.