

# CLIP-Guided Federated Learning on Heterogeneous and Long-Tailed Data

Jiangming Shi<sup>1</sup>, Shanshan Zheng<sup>2</sup>, Xiangbo Yin<sup>2</sup>, Yang Lu<sup>2</sup>, Yuan Xie<sup>3, 4\*</sup>, Yanyun Qu<sup>1, 2\*</sup>

<sup>1</sup> Institute of Artificial Intelligence, Xiamen University

<sup>2</sup> School of Informatics, Xiamen University

<sup>3</sup> East China Normal University

<sup>4</sup> Chongqing Institute of East China Normal University

jiangming.shi@outlook.com, shanshanzheng@stu.xmu.edu.cn, S\_yinx@163.com, luyang@xmu.edu.cn, yxie@cs.ecnu.edu.cn, yyqu@xmu.edu.cn

## Abstract

Federated learning (FL) provides a decentralized machine learning paradigm where a server collaborates with a group of clients to learn a global model without accessing the clients' data. User heterogeneity is a significant challenge for FL, which together with the class-distribution imbalance further enhances the difficulty of FL. Great progress has been made in large vision-language models, such as Contrastive Language-Image Pre-training (CLIP), which paves a new way for image classification and object recognition. Inspired by the success of CLIP on few-shot and zero-shot learning, we use CLIP to optimize the federated learning between server and client models under its vision-language supervision. It is promising to mitigate the user heterogeneity and class-distribution balance due to the powerful cross-modality representation and rich open-vocabulary prior knowledge. In this paper, we propose the CLIP-guided FL (CLIP2FL) method on heterogeneous and long-tailed data. In CLIP2FL, the knowledge of the off-the-shelf CLIP model is transferred to the client-server models, and a bridge is built between the client and server. Specifically, for client-side learning, knowledge distillation is conducted between client models and CLIP to improve the ability of client-side feature representation. For server-side learning, in order to mitigate the heterogeneity and class-distribution imbalance, we generate federated features to retrain the server model. A prototype contrastive learning with the supervision of the text encoder of CLIP is introduced to generate federated features depending on the client-side gradients, and they are used to retrain a balanced server classifier. Extensive experimental results on several benchmarks demonstrate that CLIP2FL achieves impressive performance and effectively deals with data heterogeneity and long-tail distribution. The code is available at <https://github.com/shijiangming1/CLIP2FL>.

## Introduction

With the demand for data privacy protection, federated learning (Shokri and Shmatikov 2015; Huang et al. 2023c) (FL) springs up, which only uses the local client models without their private data to train the server model, and it provides a distributed machine-learning paradigm with the communication efficiency and privacy protection.

\*Corresponding author.

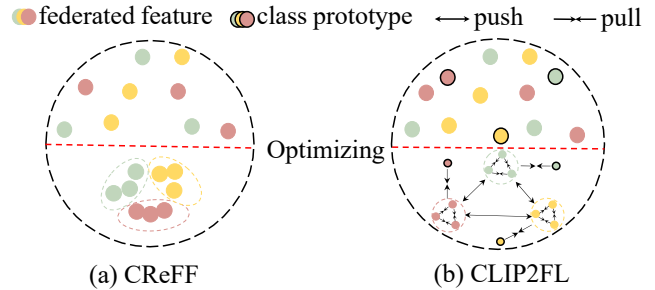


Figure 1: The figure illustrates the differences between CReFF and CLIP2FL: CReFF uses only gradients for supervision, while CLIP2FL introduces the textual features of the CLIP as prototypes for federated feature learning.

It allows users to collectively reap the benefits of a server model trained from client-side privacy data without accessing the client-side data (Li et al. 2021; Huang et al. 2023a).

Accompanied by the prospect of FL, it faces the challenge of data heterogeneity, which is a universal issue arising from disparities in data distribution among clients due to data source differences between clients (Karimireddy et al. 2020; Huang, Ye, and Du 2022). Furthermore, real-world data frequently exhibits the long-tailed phenomenon (Shuai et al. 2022), that is, the class distribution is imbalanced, where the head classes have a large number of samples while the tail classes have a small number of samples. Long-tailed data makes client-side models biased toward the head class (Shang et al. 2022). It is known that the server model based on client models is destined to not work well if the client models are poor in performance. Thus, heterogeneity and long-tailed data make FL more challenging.

There are some FL methods for imbalanced data. In the works (Duan et al. 2020; Yang et al. 2021), the class distribution is appended to adopt client selection to match clients. Because clients need to upload the local class distribution to the server, the operation violates the principle of privacy protection of FL. In FedNova (Wang et al. 2021), the auxiliary balanced data on the server are used to compute the class distribution output by the global class, but they are unavailable in the real world. CReFF (Shang et al. 2022) is the latest solution, whose core idea is to implement the decoupling

strategy to generate the class-distribution balanced federated features for the server model and to retrain the classifier by the federated features. However, in CReFF, the generation of federated features via client-side gradient information brings two limitations: 1) It is a one-to-many mapping between gradients and samples, which results in the problem becoming ill-posed; 2) It lacks semantic supervision, which could result in the federated features lacking discriminative ability for their respective classes. As shown in Figure 1, we give the comparison of CReFF and CLIP2FL in designation.

In short, the essential challenges of FL are the distinctiveness of feature representation and effective semantic supervision, but the existing methods have limitations in the solution of the two problems. Recently, large visual-language models with open-vocabulary supervision have made great progress, such as Contrastive Vision-Language Pre-training (CLIP). It paves a new way for image classification and object recognition (Chen et al. 2023). CLIP is trained on large-scale free-available image-text pairs from websites and built the vision-language correlation to achieve promising open-vocabulary recognition. It is proven that CLIP benefits the few-shot and zero-shot learning (Shen et al. 2022). However, the potential of employing CLIP to enhance learning in FL remains unexplored until now.

Considering that CLIP has rich prior knowledge and powerful vision-language cross-modality representation, we discuss how to use CLIP to solve the FL on heterogeneous and long-tailed data. There are two important issues: 1) How to use CLIP to improve the feature representation of client models; 2) How to use CLIP to mitigate the influence of heterogeneity and class-distribution imbalance on the server model. For the first issue, CLIP is used as “Teacher” while the client model is treated as “Student”, and knowledge is transferred from Teacher to Student for improving the ability of feature representation. For the second issue, inspired by the success of the retraining strategy on class-distribution-imbalance visual classification, we generate the federated features and use them to retrain the balance server model. We use CLIP to constrain the generation of federated features from the client-side gradients under its effective semantic supervision, so the generated federated features may avoid heterogeneity and class-distribution imbalance to some degree. Moreover, the federate features are related to explicit semantic cues.

Therefore, we propose the CLIP-guided federated learning method named CLIP2FL for heterogeneous and long-tailed data. CLIP2FL bridges client-side training with server-side training in FL via semantic supervision of CLIP. Firstly, in the client side, we conduct knowledge distillation to transfer the prior knowledge of an off-the-shelf CLIP model to local client models. The image and text encoders of CLIP are treated as the teacher model (Teacher) and the local client models are treated as the student model (Student). The class distribution output by “Student” is enforced to be consistent with the counterpart by CLIP. CLIP2FL significantly mitigates the bias of client-side classifiers to the head classes. Secondly, contrastive learning is implemented to generate the federated features for the server model, which are used for retraining the server classifier. The text encoder

of CLIP presents the prototype feature of a class, and the generated federated features are enforced to mimic the output features of the text encoder of CLIP.

To summarize, our contributions are threefold:

- We propose the CLIP-guided federated learning method (CLIP2FL) for heterogeneous and long-tailed data, which constructs a bridge between client and server training with CLIP supervision. To the best of our knowledge, we are the first to explore how CLIP benefits FL.
- In the client side, we conduct knowledge distillation that transfers the vision-language prior knowledge of CLIP to local client models to improve the feature representation of client models. In the server side, unlike CReFF which generates the federated features only depending on the gradient information, we employ a prototype contrastive learning to make the federated features similar to the output features of the CLIP text encoder without accessing client data. Retraining the server model on federated features effectively mitigates the heterogeneity and class-distribution imbalance.
- Extensive experimental results demonstrate that CLIP2FL is effective and efficient and it is superior to the state-of-the-art FL methods on multiple datasets.

## Related Work

### Federated Learning with Heterogeneous Data

The problem of data heterogeneity in FL is mainly caused by the fact that the data of the clients participating in the training are independently distributed and their sample distributions are not consistent, which leads to a serious decline in model accuracy. At present, there are many methods to solve this problem. They are divided into three categories: jointly improving the client and server models, stabilizing the local models, and improving the server model. The typical first category methods (Li et al. 2023; Huang et al. 2021, 2023b; Feng et al. 2023b) mainly apply optimization strategies to improve the optimization of the client and server. In the second category methods (Zhang et al. 2021; Chen et al. 2020; Zhang et al. 2022; Huang et al. 2022; Feng et al. 2023a), knowledge transfer is used to transfer the local knowledge in a model-agnostic to solve data heterogeneity. The third category methods (Lin et al. 2020; Chen and Chao 2021) aim at eliminating the differences of heterogeneous data by improving model aggregation on the server. Although the existing methods mitigate the data heterogeneity to a certain extent, they do not consider the FL on long-tailed data.

### Contrastive Language-Image Pre-training

In recent years, Vision-Language models have achieved advances, and Contrastive Language-Image Pre-training (CLIP) is a typical model. CLIP (Radford et al. 2021) is pre-trained on web-scale image-text data and shows unprecedented generality. CLIP is applied to improve many visual tasks, such as semantic segmentation (LSeg (Li et al. 2022a), GroupViT (Xu et al. 2022)), object detection (ViLD (Gu et al. 2022), GLIP (Li et al. 2022b)), image generation (CLIPasso (Vinker et al. 2022)), image retrieval (Context-12W (Tang et al. 2023a), AlignCMSS (Tang et al. 2023b))

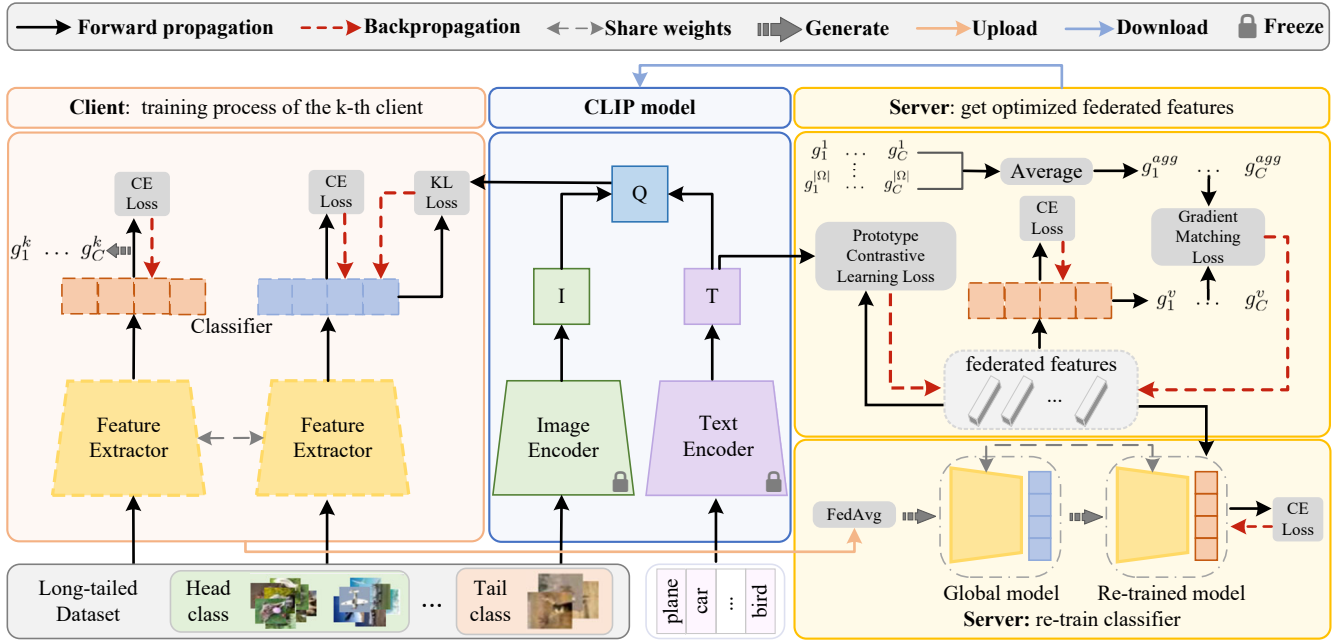


Figure 2: The framework of CLIP2FL. It includes two core components: local training on clients and classifier re-training on the server. A prior knowledge-rich CLIP model acts as a bridge to connect the two components and helps the two core components to learn better.

and video understanding (ActionCLIP (Wang, Xing, and Liu 2021), CLIP4Clip (Luo et al. 2022)). It is recognized that CLIP benefits for unseen-class classification and data deficiency (Rombach et al. 2022; Tewel et al. 2022). However, it is unexplored how CLIP benefits federated learning. Taking inspiration from the accomplishments of CLIP in few-shot and zero-shot learning, we delve into how CLIP improves FL on heterogeneous and long-tailed data.

## Method

### Problem Formulation and Notation

We aim to learn a global server model for image classification on heterogeneous and long-tailed data. CLIP2FL includes a central server model and  $K$  client models. The local-client datasets are denoted by  $\mathcal{D}^1, \mathcal{D}^2, \dots, \mathcal{D}^K$ , and the merge of local-client datasets is denoted by  $\mathcal{D} = \cup_k \mathcal{D}^k$ , where  $k \in \{1, 2, \dots, K\}$ . In this paper, the dataset  $\mathcal{D}$  comprises  $C$  classes and the class distribution  $\chi$  of  $\mathcal{D}$  is long-tailed. Let  $\mathcal{D}_c^k = \{(x, y) \in \mathcal{D}^k : y = c\}$  denotes the set of samples with ground-truth label  $c$  on the  $k$ -th client and contains  $n_c^k$  samples, so the number of samples in the  $c$ -th class is  $n_c = \sum_{k=1}^K n_c^k$ . We rearrange them according to the number of their samples in descending order as  $n_1 > \dots > n_c > \dots > n_C$ . In CLIP2FL, both the local client model and server model are classification models and contain two main components: a feature extractor  $f_\theta$  w.r.t. the parameter  $\theta$ , and a linear classifier  $g_\varphi$  w.r.t. the parameter  $\varphi$ . An input sample  $x$  is fed into a classification model, and it is processed by a feature extractor that outputs a  $d$ -dimension feature represented as  $z = f_\theta(x)$  and then it is processed by a classifier

that outputs a logit vector  $p = g_\varphi(z)$  as the label prediction for  $x$ . Here we define the classification network with the parameter  $w = \{\theta, \varphi\}$ .

### Framework of CLIP2FL

The common streamline of FL contains three steps in communication iteration: 1) The server broadcasts the global model to clients; 2) The client-side models are updated by training on their local data; 3) The client-side updated models are aggregated to the server, based on which a new global model is updated, and repeated until convergence.

Following the three steps of FL, CLIP2FL focuses on the latter two steps: the client-side learning and the server-side learning, as shown in Figure 2. For the client-side learning, we employ knowledge distillation to transfer the knowledge from CLIP to the local client models, enhancing the ability of feature representation of the client-side models. Thanks to the inherent vision-language correlation within CLIP, the problem of strong bias is alleviated after conducting knowledge distillation. As a result, the gradients of client models become more distinct and benefit the effectiveness of the ensuing server-side learning. For the server-side learning, CLIP2FL first generates federated features, which are used to retrain a balanced classifier of the server model. The class distributions of client-side private data are very different, e.g. some private data are long-tailed, and some private data only contain part of the classes, so it is impossible for the classifier of the server model to work well on heterogeneous and long-tailed data only with one-time iteration. It is necessary to retrain the classifier of the server model. The generated federated features make the server-side model less

biased to the head classes based on the retraining strategy.

### Client-side Learning

In this subsection, we introduce how CLIP benefits the local client models. Concretely, CLIP is treated as ‘‘Teacher’’, while a local client model is treated as ‘‘Student’’. CLIP offers guidance to the client models during each iteration, facilitating its updates and enhancements. We directly use a ready-made pre-training CLIP while not training it again. CLIP typically contains an image encoder  $I_{enc}$  whose input is an image and a text encoder  $T_{enc}$  whose input is the text. Given an input image  $x$  with its label name  $l$ , we can get the visual feature  $f'_i = I_{enc}(x) \in R^{d_i}$  and the text feature  $f'_r = T_{enc}(l) \in R^{d_r}$ . The class label is transformed into a sentence, e.g. ‘‘cat’’ is transformed into ‘‘This is a cat’’.

In order to mitigate the heterogeneous and long-tailed problem, the output of client models is enforced to be consistent with the output of CLIP. The loss of local training is formulated as:

$$L_{loc} = L_{ce}(y, p_k^t) + \beta \cdot KL(q_k^t || p_k^t), \quad (1)$$

where  $L_{ce}(\cdot, \cdot)$  is the cross-entropy loss,  $KL$  is the Kullback-Leibler divergence (Joyce 2011) and  $\beta$  is a hyper-parameter balanced the cross-entropy loss and the loss of knowledge distillation. The latter loss helps to narrow the output class prediction of the client and CLIP,  $p_k^t$  and  $q_k^t$  denote the output logits vector of the client  $k$  and CLIP on local data  $\mathcal{D}^k$  in the  $t$ th round, respectively. Then the  $k$ th client updates the received model similar to FedAvg (McMahan et al. 2017) as follows:

$$w_k^{t+1} \leftarrow w_k^t - \eta \nabla_w L_{loc}(w^t; \mathcal{D}^k). \quad (2)$$

### Server-side Learning

In this subsection, we introduce the server-side learning. Server-side learning comprises three important components: aggregation of the gradients of the client-side classifiers corresponding to their client-side models  $\hat{\varphi}^t$ , the generation of federated features, and retraining the global model.

**Gradient Aggregation.** We compute the gradients of the current server classifier  $\hat{\varphi}^t$ , which is used to optimize the federated feature for re-training a balanced classifier on the server. Specifically, the current server classifier  $\hat{\varphi}^t$  replaces the classifier of the  $k$ th client model, and with the resampled data from the dataset  $\mathcal{D}_i$ , the new client model produces  $d$ -dimension real sample features  $\mathcal{Z}_c^k = \{z_{c,i}^k\}_{i=1}^{n_c^k}$  for the  $c$ th class. Then we use the server classifier  $\hat{\varphi}^t$  to compute the gradients  $g_c^k \in R^{C \times d}$  of the  $c$ th class as below:

$$g_c^k = \frac{1}{n_c^k} \sum_{i=1}^{n_c^k} \nabla_{\hat{\varphi}^t} L_{ce}(z_{c,i}^k, y_i). \quad (3)$$

After that, the real gradients  $\{g_c^k | c \in C^k\}$  are uploaded to the server. Due to data heterogeneity,  $C^k$  here indicates that only the class index set is involved in the current client training data. Importantly, CLIP2FL only obtains gradients subsequent to the irreversible averaging operation, without

any direct access to the client data.

**Federated Feature Generation.** Firstly, the server aggregates real gradients of each class and computes the average of the real gradients over all selected clients in  $\Omega^t$ :

$$g_c^{agg} = \frac{1}{|\Omega_c^t|} \sum_{k=1}^{|\Omega_c^t|} g_c^k, \quad (4)$$

where  $\Omega_c^t$  is the subset of clients including the  $c$ th class.

After that, we generate a set of  $d$ -dimension federated features of size  $m$  for each class in the  $t$ -th round, denoted as  $V_c^t = \{v_{c,i}^t\}_{i=1}^m$ . The federated features are generated by enforcing the two constraints: 1) The gradient computed by the server classifier  $\hat{\varphi}^t$  on the federated features should be consistent with the gradient computed by the real features; 2) The federated features should be similar to the class prototypes that are produced by the text encoder of CLIP. Due to highly distinctive class representations, we consider that prototypes are helpful to mitigate the strong bias problem.

As for the first constraint, we use the classifier  $\hat{\varphi}^t$  to generate the gradient of the corresponding class for the federated features:

$$g_c^v = \frac{1}{m} \sum_{i=1}^m \nabla_{\hat{\varphi}^t} L_{ce}(v_{c,i}^t, y_i). \quad (5)$$

To further ensure consistency between federated features and real features, the gradient matching loss (Zhao, Mopuri, and Bilen 2021) is used to measure the difference between the gradients generated by the server classifier for the two types of features and minimize it:

$$L_{grad} = D(g_c^v, g_c^{agg}) = \frac{1}{C} \sum_{j=1}^C \left( 1 - \frac{g_c^v[j] \cdot g_c^{agg}[j]}{\|g_c^v[j]\| \times \|g_c^{agg}[j]\|} \right), \quad (6)$$

where  $g[j]$  denotes the  $j$ th row of the gradient. This helps to make the federated features as similar as possible to the real features.

As for the second constraint, we use the prototype output by the text encoder of CLIP with strong semantic information to guide the generation of federated features. A prototype with strong supervision information benefits the generation of distinctive federated features. Concretely, we conduct prototype contrastive learning. The federated features  $V_c^t$  and their intra-class prototype  $f_r^c$  are treated as positive pair-wise samples while federated features  $V_c^t$  and their inter-class federated features are treated as negative pair-wise samples. The loss of prototype contrastive learning is:

$$L_{pcl} = \sum_{i=1}^{C \times m} -\log \frac{\exp(\langle v_{c,i}, f_r^c \rangle / \tau)}{\sum_{j=1}^{C \times m} 1_{[j \neq i]} \exp(\langle v_{c,i}, v_j \rangle / \tau)}, \quad (7)$$

where  $\langle \cdot, \cdot \rangle$  denotes cosine similarity and  $\tau$  is a temperature hyper-parameter. The total loss for obtaining optimized federated features is as follows:

$$L_{total} = L_{grad} + \eta \cdot L_{pcl}, \quad (8)$$

Type	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+ $\tau$ -norm	49.95	51.41	72.08	26.22	33.71	43.65
SOTA	CReFF	70.55	73.08	80.71	34.67	37.64	47.08
Proposed method	CLIP2FL	<b>73.37</b> ( $\uparrow$ 2.82)	<b>75.35</b> ( $\uparrow$ 2.27)	<b>81.18</b> ( $\uparrow$ 0.47)	<b>37.56</b> ( $\uparrow$ 2.89)	<b>41.29</b> ( $\uparrow$ 3.65)	<b>48.20</b> ( $\uparrow$ 1.12)

Table 1: Top-1 classification accuracy(%) on CIFAR-10-LT and CIFAR-100-LT datasets with different FL methods, where the results are referred in (Shang et al. 2022). The best results are marked in bold.

where  $\eta$  is a hyper-parameter controlling the loss weights.

After the federated features  $V^t$  have been optimized using the total loss function, we obtain an updated  $V^{t+1}$  for retraining the classifier. Applying the two constraints to the generation of federated features provides two significant advantages. On the one hand, CReFF (Shang et al. 2022) only uses gradient information, which is not enough to guide the generation of federated features (Eq. 6). Due to privacy protection, the gradient generated by the client is uploaded after an irreversible average operation, so the averaged gradient information and real features are one-to-many, which means the federated features are not very distinguishable, which is not conducive to subsequent classifier re-training. On the other hand, we utilize the CLIP model, which has rich prior knowledge, as an auxiliary module to assist in the training of federated features on the server. This allows the use of the rich semantic information contained within the CLIP model to guide the generation of federated features, resulting in more accurate and useful federated features. This neatly ties into the training of the client and shows how to make good use of vision-language models for FL.

**Classifier Re-Training.** After the local client model is updated, in the  $t$ th round, clients in the selected set  $\Omega^t$  upload their updated models  $w_k^{t+1}$  to the server. The average of parameters of selected client models is assigned to the global model. The server model performs a weighted average to update the global model for the  $t+1$  round where the weight of a class is its proportion to the amount of the whole dataset:

$$w^{t+1} = \sum_{k \in \Omega^t} \frac{|\mathcal{D}^k|}{\sum_{k \in \Omega^t} |\mathcal{D}^k|} w_k^{t+1}. \quad (9)$$

The retraining model is the copy of the global model, which is retrained by fixing the feature extractor, while the server classifier  $g_\varphi$  w.r.t  $\hat{\varphi}^t$  is updated with federated features and get the new parameter  $\hat{\varphi}^{t+1}$  for round  $t+1$ .

$$\hat{\varphi}^{t+1} \leftarrow \hat{\varphi}^t - \eta \nabla_{\hat{\varphi}} L_{ce}(v_i^t, y_i). \quad (10)$$

Finally, we obtain the re-trained model  $\hat{w}^{t+1} = \{\theta^{t+1}, \hat{\varphi}^{t+1}\}$  and the updated global model  $w^{t+1} =$

---

#### Algorithm 1: Training Process for Round $t$

---

**Input:** global model parameter  $w^t = \{\theta^t, \varphi^t\}$  and current re-trained classifier  $\hat{\varphi}^t$ , randomly initialized federated features  $\{V_{c,i}^t \mid c \in C^k\}_{i=1}^m$

**Output:** a re-trained model parameterized as  $\hat{w}^{t+1} = \{\theta^{t+1}, \hat{\varphi}^{t+1}\}$  for next round

- 1: # Server executes:
  - 2: Send  $w^t$  to clients;
  - 3: Randomly select a set of online clients  $\Omega^t$ ;
  - 4: # Client executes:
  - 5: **for**  $k \in \Omega^t$  **do**
  - 6:   Update local model  $w_k^{t+1}$  by Eq.1 and Eq.2;
  - 7:   compute real feature gradients  $\{g_c^k \mid c \in C^k\}$  using Eq.3;
  - 8: **end for**
  - 9: Send  $\{w_k^{t+1}, g_c^k \mid k \in \Omega^t\}$  to the server;
  - 10: # Server executes:
  - 11: Update the global model  $w^{t+1}$  by Eq.9;
  - 12: Aggregate real feature gradients to  $g_c^{agg}$  using Eq.4;
  - 13: Compute federated feature gradients  $g_c^v$  by Eq.5
  - 14: Optimize federated features to  $V^{t+1}$  by gradient matching loss and prototype contrastive learning loss (Eq.6, Eq.7 and Eq.8)
  - 15: Re-train the classifier  $\hat{\varphi}^t$  to  $\hat{\varphi}^{t+1}$
  - 16: Send  $w^{t+1} = \{\theta^{t+1}, \varphi^{t+1}\}$  and  $\hat{w}^{t+1} = \{\theta^{t+1}, \hat{\varphi}^{t+1}\}$  to clients.
- 

$\{\theta^{t+1}, \varphi^{t+1}\}$ . These models are then sent to clients for the next round. The comprehensive training process is presented in Algorithm 1.

## Experiment

### Experiment Setup

**Datasets.** We implement CLIP2FL on three frequently used datasets with the long-tailed data: CIFAR-10/100-LT (Krizhevsky, Hinton et al. 2009) and ImageNet-LT (Rus-

Type	Method	ImageNet-LT			
		All	Many	Medium	Few
Heterogeneity-oriented FL methods	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
Imbalance-oriented FL methods	Fed-Focal Loss	21.60	31.74	15.77	5.52
	Ratio Loss	24.31	36.33	18.14	7.41
	FedAvg+ $\tau$ -norm	21.58	31.66	15.76	4.33
SOTA	CReFF	26.31	<b>37.44</b>	21.87	10.29
Proposed method	CLIP2FL	<b>27.72</b> ( $\uparrow$ 1.41)	35.06 ( $\downarrow$ 2.38)	<b>27.55</b> ( $\uparrow$ 5.68)	<b>24.80</b> ( $\uparrow$ 14.51)

Table 2: Top-1 classification accuracy(%) on ImageNet-LT dataset with different FL methods, where the results are referred in (Shang et al. 2022). The best results are marked in bold.

sakovsky et al. 2015). As for the first two datasets, we follow the previous work (Cui et al. 2019) to sample these two datasets with different imbalance factors (IF = 100, 50, 10) for long-tailed distribution, and we follow CReFF (Shang et al. 2022) to use Dirichlet distribution with the key parameter  $\alpha$  to generate the heterogeneous data partition among clients, where the value of  $\alpha$  is set to 0.5 on CIFAR-10/100-LT. ImageNet-LT (Russakovsky et al. 2015) has 115.8K images from 1000 classes and the number of images per class ranging from 1280 to 5, where the value of  $\alpha$  is set to 0.1.

**Implementation.** ResNet-8 is used as the feature extractor on CIFAR-10/100-LT and ResNet-50 is used for ImageNet-LT. Note that in order to work with CLIP, a layer of MLP is added to both ResNet-8 and ResNet-50 to make their features have the same dimension as the output of CLIP. CLIP with the ViT-B/32 model is used in our experiments.

The number of clients is set to 20, and 40% of them are randomly selected as online clients to participate in training. The batch size of client-side training is set to 32 for all datasets and we set the number of federated features to 100 for each class. Experiments were conducted using PyTorch on four NVIDIA GeForce RTX 3090 GPUs. We employed the standard cross-entropy loss by default and executed 200 communication rounds.

## Experiment Results

We compare CLIP2FL with eleven FL methods: FedAvg (McMahan et al. 2017), FedAvgM (Hsu, Qi, and Brown 2019), FedProx (Li et al. 2020), FedDF (Lin et al. 2020), FedBE (Chen and Chao 2021), CCVR (Luo et al. 2021) and FedNova (Wang et al. 2020), Fed-Focal Loss (Sarkar, Narang, and Rai 2020), Ratio Loss (Wang et al. 2021) and FedAvg with  $\tau$ -norm (Kang et al. 2020), CReFF (Shang et al. 2022). The first seven methods are heterogeneity-oriented methods, and Fed-Focal Loss (Sarkar, Narang, and Rai 2020), Ratio Loss (Wang et al. 2021) and FedAvg with  $\tau$ -norm (Kang et al. 2020) are imbalance-oriented methods.

**Results on CIFAR-10/100-LT.** In Table 1, we compare the classification accuracy of our CLIP2FL with various FL methods on CIFAR-10-LT and CIFAR-100-LT. It is ob-

served that CLIP2FL achieves the best classification accuracy on both CIFAR-10-LT and CIFAR-100-LT. When IF = 50, CLIP2FL gains 2.27% on CIFAR-10-LT and 3.65% on CIFAR-100-LT in classification accuracy compared with CReFF, respectively. Even when IF = 100 which causes the data to be severely imbalanced, CLIP2FL achieves the gain by nearly 3%. It fully demonstrates the superiority and effectiveness of CLIP2FL.

**Results on ImageNet-LT.** In Table 2, we show the classification accuracy comparison of CLIP2FL with different FL methods on ImageNet-LT. As shown in Table 2, the four divisions are defined as follows: “Many” (> 100 samples), “Medium” (20 – 100 samples), “Few” (< 20 samples), and “All” (overall accuracy), respectively. It is observed that CLIP2FL is superior to the compared methods in the accuracy of the above four divisions, especially CLIP2FL achieves a significant gain of 5.68% and 14.51% on the “Medium” and “Few” divisions compared with CReFF. It indicates that CLIP2FL not only enhances the overall performance but also obviously improves the classification performance of the tail categories. The comparison on ImageNet-LT further highlights the advantages of CLIP2FL, especially when dealing with long-tailed datasets.

## Model Validation

**Ablation Study.** We conduct ablation studies with two losses on CIFAR-10/100-LT: 1) The knowledge distillation loss from CLIP to the client-side model (defined as  $KL$ ); 2) The prototype contrastive learning loss on the server side for generating federated features (defined as  $L_{pcl}$ ). We select the method (Shang et al. 2022) as our baseline. As shown in Table 3, each component is revealed and plays an important role in our CLIP2FL. CLIP2FL only with  $KL$  is superior to the baseline method with the gain of 1.65% and 1.21% on CIFAR-10/100-LT, which shows CLIP can help the client-side model to improve the ability of feature representation. At the same time, the baseline method with the two losses achieves the gain of 2.82% and 2.89% on CIFAR-10/100-LT. It shows that CLIP benefits to supervise the generation of federated features. In addition, this result shows that  $KL$  and  $L_{pcl}$  can complement each other.

Method	$KL$	$L_{pcl}$	CIFAR-10-LT	CIFAR-100-LT
Baseline			70.55	34.67
w/o $L_{pcl}$	✓		72.20	35.88
CLIP2FL	✓	✓	<b>73.37</b>	<b>37.56</b>

Table 3: Ablation study to investigate the effectiveness of CLIP2FL on CIFAR-10-LT and CIFAR-100-LT with  $IF = 100$ .




Figure 3: Visualization of CKA similarity at the classifier layer under the “client model-client model” setting with or without using the CLIP model to guide the client model.

**Visualization.** In order to further verify that CLIP benefits to guide the client learning and to alleviate classifier bias caused by long-tailed data, we follow the CCVR (Luo et al. 2021) and adopt Centered Kernel Alignment (CKA) (Kornblith et al. 2019) to measure the similarity of the output features between two local client models given the same input. CKA outputs a similarity score between 0 (not similar at all) and 1 (identical). The visualization results are shown in Figure 3, the darker color indicates a higher degree of similarity. Compared with the w/o CLIP model (right), the CLIP-guided client model (left) significantly improves the CKA similarity between different client models, it demonstrates that using the CLIP model to guide client learning can effectively reduce the bias of the classifier, which means improving feature representation learning more fully.

**Hyperparameter Analysis.** Three important hyperparameters in our CLIP2FL are  $\beta$ ,  $\eta$  and  $m$ .  $\beta$  controls the distillation weight of the CLIP model to guide the client model’s training,  $\eta$  controls the balance of prototype contrastive learning loss and gradient matching loss when optimizing federated features, and  $m$  is the number of federated features initialized for each class. We analyzed hyperparameters under the condition that only one hyperparameter is selected as the variable, while all other hyperparameters are kept constant. We found that CLIP2FL achieved the best performance when  $\beta = 3.0$ ,  $\eta \in \{0.001, 0.0001, 1e-5\}$  and  $m = 100$ . Surprisingly, we observed that the performance of CLIP2FL exhibits limited sensitivity to the hyperparameters, indicating a degree of robustness in the CLIP2FL. More details are in the Appendix.

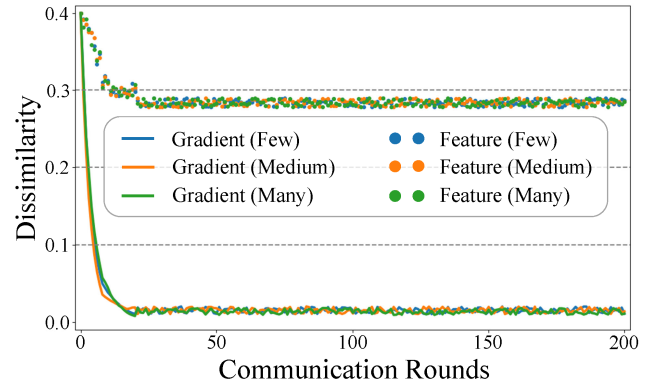


Figure 4: Analysis of two dissimilarities on CIFAR-10-LT with  $IF = 100$ . The dotted lines denote the dissimilarities between federated features and real features, and the solid lines denote the dissimilarities between federated feature gradients and real feature gradients.

### Dissimilarity Analysis of Gradient & Feature

We discuss how the dissimilarity between the federated features gradients and the real sample features gradients change with the increase of iterations (Eq. 8). Here, the classes of CIFAR-10-LT are divided into three groups: Many ( $> 1500$  samples), Medium ( $200 - 1500$  samples) and Few ( $< 200$  samples), and the average dissimilarities for each group are calculated separately. As the solid lines shown in Figure 4, after several rounds of communication, the dissimilarities between the two gradients tend to an approximate value of 0. It shows that the loss can successfully optimize the gradients generated by the federated features to be consistent with the gradients of the real sample features. It indicates that the re-training classifier with federated features on the server side is effective. Moreover, we also calculate cosine dissimilarity between the federated features and the real sample features, which is shown in the dotted lines in Figure 4. With the increase of training rounds, the dissimilarities between federated features and real sample features are also decreasing but are larger than 0. It demonstrates that although the corresponding gradients are very close, the generated federated features are not completely similar to the real sample features, which ensures the privacy of real sample features.

### Conclusion

In this work, we provide the CLIP-guided FL method (CLIP2FL) for federated learning on heterogeneity and long-tailed data. On the client side, a ready-made CLIP model is used for knowledge distillation to the client models. On the server side, CLIP is used to guide the generation of federated features via prototype contrastive learning. The federated features are generated to mitigate the strong bias problem caused by long-tailed data. CLIP2FL enforces the semantic supervision on FL and bridges the client-side and server-side training with CLIP. The experimental results on three benchmarks achieve a significant superiority against the state-of-the-art methods.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grants (62176224, 62222602, 62176092, 62376233), Natural Science Foundation of Chongqing (CSTB2023NSCOJ0X0007), CCF-Lenovo Blue Ocean Research Fund.

## References

- Chen, H.; and Chao, W. 2021. FedBE: Making Bayesian Model Ensemble Applicable to Federated Learning. In *ICLR*.
- Chen, R.; Liu, Y.; Kong, L.; Zhu, X.; Ma, Y.; Li, Y.; Hou, Y.; Qiao, Y.; and Wang, W. 2023. CLIP2Scene: Towards Label-efficient 3D Scene Understanding by CLIP. *CoRR*, abs/2301.04926.
- Chen, Y.; Qin, X.; Wang, J.; Yu, C.; and Gao, W. 2020. FedHealth: A Federated Transfer Learning Framework for Wearable Healthcare. *IEEE Intell. Syst.*, 83–93.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *CVPR*, 9268–9277.
- Duan, M.; Liu, D.; Chen, X.; Liu, R.; Tan, Y.; and Liang, L. 2020. Self-balancing federated learning with global imbalanced data in mobile systems. *IEEE Transactions on Parallel and Distributed Systems*, 59–71.
- Feng, C.-M.; Li, B.; Xu, X.; Liu, Y.; Fu, H.; and Zuo, W. 2023a. Learning Federated Visual Prompt in Null Space for MRI Reconstruction. In *CVPR*, 8064–8073.
- Feng, C.-M.; Yu, K.; Liu, N.; Xu, X.; Khan, S.; and Zuo, W. 2023b. Towards Instance-adaptive Inference for Federated Learning. In *ICCV*, 23287–23296.
- Gu, X.; Lin, T.; Kuo, W.; and Cui, Y. 2022. Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. In *ICLR*.
- Hsu, T. H.; Qi, H.; and Brown, M. 2019. Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification. *CoRR*, abs/1909.06335.
- Huang, H.; Shang, F.; Liu, Y.; and Liu, H. 2021. Behavior Mimics Distribution: Combining Individual and Group Behaviors for Federated Learning. In *IJCAI*, 2556–2562.
- Huang, W.; Ye, M.; and Du, B. 2022. Learn from others and be yourself in heterogeneous federated learning. In *CVPR*, 10143–10153.
- Huang, W.; Ye, M.; Du, B.; and Gao, X. 2022. Few-Shot Model Agnostic Federated Learning. In *ACMMM*, 7309–7316.
- Huang, W.; Ye, M.; Shi, Z.; and Du, B. 2023a. Generalizable Heterogeneous Federated Cross-Correlation and Instance Similarity Learning. *TPAMI*.
- Huang, W.; Ye, M.; Shi, Z.; Li, H.; and Du, B. 2023b. Rethinking federated learning with domain shift: A prototype view. In *CVPR*, 16312–16322.
- Huang, W.; Ye, M.; Shi, Z.; Wan, G.; Li, H.; Du, B.; and Yang, Q. 2023c. A Federated Learning for Generalization, Robustness, Fairness: A Survey and Benchmark. *arXiv*.
- Joyce, J. M. 2011. Kullback-Leibler Divergence. In *International Encyclopedia of Statistical Science*, 720–722.
- Kang, B.; Xie, S.; Rohrbach, M.; Yan, Z.; Gordo, A.; Feng, J.; and Kalantidis, Y. 2020. Decoupling Representation and Classifier for Long-Tailed Recognition. In *ICLR*.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S. J.; Stich, S. U.; and Suresh, A. T. 2020. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. In *ICML*, 5132–5143.
- Kornblith, S.; Norouzi, M.; Lee, H.; and Hinton, G. 2019. Similarity of neural network representations revisited. In *ICML*, 3519–3529.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, B.; Schmidt, M. N.; Alstrøm, T. S.; and Stich, S. U. 2023. On the effectiveness of partial variance reduction in federated learning with heterogeneous data. In *CVPR*, 3964–3973.
- Li, B.; Weinberger, K. Q.; Belongie, S. J.; Koltun, V.; and Ranftl, R. 2022a. Language-driven Semantic Segmentation. In *ICLR*.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.; Chang, K.; and Gao, J. 2022b. Grounded Language-Image Pre-training. In *CVPR*, 10955–10965.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *MLSys*, 429–450.
- Li, X.; Jiang, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021. FedBN: Federated Learning on Non-IID Features via Local Batch Normalization. In *ICLR*.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble Distillation for Robust Model Fusion in Federated Learning. In *NeurIPS*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. CLIP4Clip: An empirical study of CLIP for end to end video clip retrieval and captioning. *Neurocomputing*, 293–304.
- Luo, M.; Chen, F.; Hu, D.; Zhang, Y.; Liang, J.; and Feng, J. 2021. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *NeurIPS*, 5972–5984.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *AISTIS*, 1273–1282.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*, 10674–10685.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 211–252.

- Sarkar, D.; Narang, A.; and Rai, S. 2020. Fed-Focal Loss for imbalanced data classification in Federated Learning. *CoRR*, abs/2011.06283.
- Shang, X.; Lu, Y.; Huang, G.; and Wang, H. 2022. Federated Learning on Heterogeneous and Long-Tailed Data via Classifier Re-Training with Federated Features. In *IJCAI*, 2218–2224.
- Shen, S.; Li, L. H.; Tan, H.; Bansal, M.; Rohrbach, A.; Chang, K.; Yao, Z.; and Keutzer, K. 2022. How Much Can CLIP Benefit Vision-and-Language Tasks? In *ICLR*.
- Shokri, R.; and Shmatikov, V. 2015. Privacy-Preserving Deep Learning. In *SIGSAC*, 1310–1321.
- Shuai, X.; Shen, Y.; Jiang, S.; Zhao, Z.; Yan, Z.; and Xing, G. 2022. BalanceFL: Addressing class imbalance in long-tail federated learning. In *IPSN*, 271–284.
- Tang, Y.; Yu, J.; Gai, K.; Jiamin, Z.; Xiong, G.; Hu, Y.; and Wu, Q. 2023a. Context-I2W: Mapping Images to Context-dependent Words for Accurate Zero-Shot Composed Image Retrieval. *arXiv preprint arXiv:2309.16137*.
- Tang, Y.; Yu, J.; Gai, K.; Wang, Y.; Hu, Y.; Xiong, G.; and Wu, Q. 2023b. Align before Search: Aligning Ads Image to Text for Accurate Cross-Modal Sponsored Search. *arXiv preprint arXiv:2309.16141*.
- Tewel, Y.; Shalev, Y.; Schwartz, I.; and Wolf, L. 2022. ZeroCap: Zero-Shot Image-to-Text Generation for Visual-Semantic Arithmetic. In *CVPR*, 17897–17907.
- Vinker, Y.; Pajouheshgar, E.; Bo, J. Y.; Bachmann, R. C.; Bermanno, A. H.; Cohen-Or, D.; Zamir, A.; and Shamir, A. 2022. CLIPasso: semantically-aware object sketching. *ACM Trans. Graph.*, 86:1–86:11.
- Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; and Poor, H. V. 2020. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization. In *NeurIPS*.
- Wang, L.; Xu, S.; Wang, X.; and Zhu, Q. 2021. Addressing class imbalance in federated learning. In *AAAI*, 10165–10173.
- Wang, M.; Xing, J.; and Liu, Y. 2021. ActionCLIP: A New Paradigm for Video Action Recognition. *CoRR*, abs/2109.08472.
- Xu, J.; Mello, S. D.; Liu, S.; Byeon, W.; Breuel, T. M.; Kautz, J.; and Wang, X. 2022. GroupViT: Semantic Segmentation Emerges from Text Supervision. In *CVPR*, 18113–18123.
- Yang, M.; Wang, X.; Zhu, H.; Wang, H.; and Qian, H. 2021. Federated learning with class imbalance reduction. In *EU-SIPCO*, 2174–2178.
- Zhang, J.; Guo, S.; Ma, X.; Wang, H.; Xu, W.; and Wu, F. 2021. Parameterized Knowledge Transfer for Personalized Federated Learning. In *NeurIPS*, 10092–10104.
- Zhang, L.; Shen, L.; Ding, L.; Tao, D.; and Duan, L. 2022. Fine-tuning Global Model via Data-Free Knowledge Distillation for Non-IID Federated Learning. In *CVPR*, 10164–10173.
- Zhao, B.; Mopuri, K. R.; and Bilal, H. 2021. Dataset Condensation with Gradient Matching. In *ICLR*.