

FedReMa: Improving Personalized Federated Learning via Leveraging the Most Relevant Clients

Han Liang^a, Ziwei Zhan^a, Weijie Liu^a, Xiaoxi Zhang^{a,*}, Chee Wei Tan^b and Xu Chen^a

^aSun Yat-sen University

^bNanyang Technological University

Abstract. Federated Learning (FL) is a distributed machine learning paradigm that achieves a globally robust model through decentralized computation and periodic model synthesis, primarily focusing on the global model's accuracy over aggregated datasets of all participating clients. Personalized Federated Learning (PFL) instead tailors exclusive models for each client, aiming to enhance the accuracy of clients' individual models on specific local data distributions. Despite of their wide adoption, existing FL and PFL works have yet to comprehensively address the class-imbalance issue, one of the most critical challenges within the realm of data heterogeneity in PFL and FL research. In this paper, we propose *FedReMa*, an efficient PFL algorithm that can tackle class-imbalance by 1) utilizing an *adaptive inter-client co-learning* approach to identify and harness different clients' expertise on different data classes throughout various phases of the training process, and 2) employing distinct aggregation methods for clients' feature extractors and classifiers, with the choices informed by the different roles and implications of these model components. Specifically, driven by our experimental findings on inter-client similarity dynamics, we develop *critical co-learning period* (CCP), wherein we introduce a module named *maximum difference segmentation* (MDS) to assess and manage task relevance by analyzing the similarities between clients' logits of their classifiers. Outside the CCP, we employ an additional scheme for model aggregation that utilizes historical records of each client's most relevant peers to further enhance the personalization stability. We demonstrate the superiority of our *FedReMa* in extensive experiments. The code is available at <https://github.com/liangh68/FedReMa>.

Keywords: Personalized Federated Learning, Class-Imbalance, Relevant Matching

1 Introduction

Federated learning (FL) enables collaborative model training across decentralized devices or data sources, preserving privacy and scalability while harnessing the collective intelligence of distributed data [18]. However, this algorithm based on periodic synchronization requires all client models to share the same structure and parameters, which brings about huge heterogeneity problems [31, 33, 30].

Unlike traditional FL approaches that view data heterogeneity across clients as an obstacle, personalized federated learning (PFL) aims to exploit these data variations and train distinct models so as to tailor clients' individual models to their local optimization subproblems under heterogeneous data distributions, focusing more

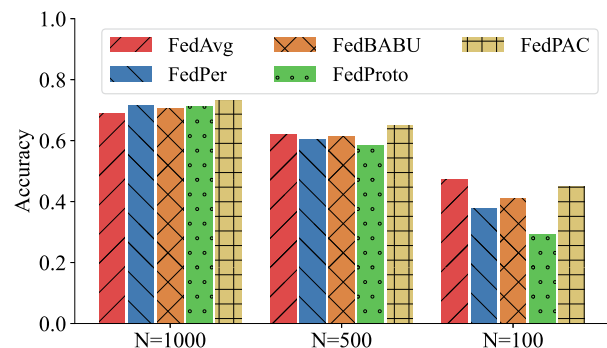


Figure 1: Average local accuracy of advanced PFL methods can even fall below FedAvg when the number of trainable samples decreases.

on model personalization. This paradigm shift in machine learning opens up new possibilities for personalized services, recommendations, and predictions, all while respecting user privacy.

Nevertheless, PFL still encounters numerous challenges. Firstly, disparate data distributions arise from varying data preferences, causing label skew and feature skew, known as the non-IID data challenge. Among all the subcategories of non-IID data issue, class-imbalance is arguably the most challenging one. While numerous studies have dedicated to solving class-imbalance in centralized machine learning settings [25, 29, 11], these methods are not readily applicable to PFL, given its unique challenges posed by the distributed and privacy-preserving nature, which prevents sharing individual client data to balance the overall dataset. Moreover, uniform aggregation operations adopted in most FL studies employed during each period of model synchronization can be detrimental from a local perspective [2]. This occurs because each model, meticulously trained for a specific client's data, may amalgamate with models harboring *unfavorable knowledge* from diverse distributions of other clients. Therefore, numerous PFL studies propose various methods to mitigate performance degradation due to standard model aggregation under non-IID data. Many of the leading PFL studies focus on retaining personalization layers locally [20, 1, 4, 2, 16] or adapting them to local tasks through certain modifications [24, 21, 8, 22, 5, 17, 23].

However, in practice, PFL may encounter scenarios where clients' data can be both sparse and imbalanced. Most of the PFL works do not fully leverage the *dynamic* inter-client relationships, which should be exploited to improve model generalization along with personalization [20, 1, 4, 22, 32]. Moreover, methods using personalized layers, represented by FedPer [1] and FedRep [4], do not enable *classifier synchronization* among clients and the PS, and they show dis-

* Corresponding author: Xiaoxi Zhang (zhangxx89@mail.sysu.edu.cn)

appointing performance. where an extensive network encompasses a multitude of clients with sparse data, PFL methods struggle significantly with generalization. Our experiments (Figure 1) demonstrate that, even the vanilla algorithm FedAvg [18] can outperform PFL methods when the number data samples decreases. Besides, the method like FedPer, which disables classifier sharing, suffer from a more obvious performance degradation than FedPAC [28], a representative method that leverages classifier collaboration. This result reveals the importance of the knowledge transfer from clients' classifiers for improving generalization in PFL. We then gain our first insight: *generalization capability needs to be retained for PFL, even though personalization is its primary goal.*

Exploiting inter-client classifier collaboration. Given the above observation, it is imperative to employ a method that can gauge the underlying correlation and leverage the valuable inter-client knowledge to optimize local classifiers. Some studies [3, 6] suggest assigning higher selection probabilities to clients with larger loss values for aggregation. However, selection based on training indicators often ignores the correlation between clients. A few works improve this strategy. FedPer++ [27] uses a feature regularization training strategy to mitigate local overfitting risks and reduce parameter differences between clients, thereby promoting global feature extractor aggregation. A more recent work, FedPAC [28], quantifies the benefit of a combination of classifiers for each client as a function of the combination weights and derives an optimization problem to estimate the optimal weights. Yet, acquiring the optimal weights in FedPAC necessitates solving a complex optimization problem, significantly prolonging training time. Moreover, this algorithm still exhibits suboptimal performance in scenarios involving sparse or homogeneous data. Therefore, existing classifier collaboration needs to be improved in terms both time efficiency and the ability to boost model generalizability and personalizability. Intuitively, these two capabilities might seem to conflict with each other. However, our second insight implies that, *selectively knowledge utilization from peers, whose models can preserve and enhance the strength in its own dominant classes while improving performance in other classes without compromising their existing strength, can improve both generalization and personalization.* An intrinsic question then arises: How to design a lightweight method to identify the most relevant peers and effectively utilize their knowledge for each individual client to achieve these capabilities? To answer this question, we make the following technical contributions:

- We introduce an adaptive approach called *FedReMa* that leverages the expertise of different clients on specific data classes throughout the training process. This approach helps address class-imbalance issues by dynamically identifying and harnessing clients' expertise in different phases.
- We identify a critical co-learning period (CCP) where we measure the task relevance using the similarity of clients' logits. To filter out irrelevant contributions, we introduce the maximum difference segmentation module (MDS). Additionally, outside the CCP, we employ a scheme based on historical clients' selections to enhance personalization stability.
- Through extensive experiments, we demonstrate the superiority of our algorithm, showcasing its effectiveness in personalized federated learning scenarios.

2 Related Work

In this section, we will present our investigation on existing FL works under non-IID data and various PFL studies. We elaborate their pros and cons, summarizing how they differ from this work.

2.1 Data Heterogeneity in Traditional FL

Data heterogeneity poses a significant challenge in federated learning, even with classical algorithms like FedAvg [18], which achieves a global consensus model through local updates and synchronization. Although FedAvg can train models efficiently with distributed data without compromising privacy, the data heterogeneity among clients presents a notable challenge. Previous studies [14] indicate that heterogeneous data not only hampers the model's generalization ability but also hinders the model's convergence.

The performance degradation can be attributed to several factors. Biased distributions, as highlighted in prior research [11, 29], are interpreted as meaningful information in machine learning. Highly imbalanced data can lead to label bias in classifiers, where decision boundaries are heavily influenced by the majority class. With inconsistent data distributions across clients, decision boundaries also vary, making it challenging to reach a consensus.

To address these issues, traditional FL algorithms have been developed based on FedAvg. FedProx [13] introduces a proximal term to the local objective function to restrict the updated models to be close to the initial global model. SCAFFOLD [10] introduces control variables to correct the client-drifts in local updates. MOON [12] takes advantage of the fact that the global model is expected to extract better features by adding a model contrastive-loss term to the loss function to learn a better feature extractor.

However, these algorithms are unable to recognize the similarity of data across different clients, which may introduce irrelevant knowledge. In contrast, our algorithm is able to exploit similarities between clients, thus improving the model's personalization capabilities.

2.2 Personalized Federated learning

Personalized federated learning (PFL) allows each client to maintain a private model to enhance personalization capability. Some work adopts the parameter decoupling approach, which treats model as the combination of a feature extractor and a classifier. To address data heterogeneity, LG-FedAvg [16] only aggregates classifiers, while FedPer [1] and FedRep [4] only aggregate the feature extractors. Further, based on FedPer, in order to prevent the model from failing to converge, FedVF [19] allows the classifier to perform low-frequency periodic aggregations, while FedPer++ [27] and FedPAC [28] employ fine-grained classifier aggregation and feature alignment to mitigate the problem of overfitting. Fed-RoD [2] allows the model to have two classifiers to satisfy both generalization and personalization performance, where the personalized classifier is trained using class-balanced loss and the generic classifier is trained using vanilla loss. FedFA [32] achieves personalization by fine-tuning classifiers with feature anchors. CReFF [21] proposes a retraining algorithm to handle the biased classifiers. FedClassAvg [8] only aggregates classifiers to seek a consensus decision boundary, and apply representation learning to stabilize the decision boundary and improve the feature extraction capabilities.

Apart from parameter decoupling, a number of machine learning techniques have also shown great potential in the field of PFL. Meta-learning based FL aims to seek an initial model that can be quickly adapted to local data with few local adaptations. Per-FedAvg [5] and pFedMe [22] focuses on optimizing the personalized models while keeping them close to the global model by introducing a regularization term to the loss function. Knowledge distillation based FL enables efficient communication and heterogeneous models by utilizing the predictive logic of the teacher model to perform local up-

dates. FedDistill [9] rectifies the local dataset by federated augmentation and performs model synchronization using global-average logits. FedKD [26] adopts adaptive mutual knowledge distillation with a shared mentor model for model synchronization and exploits dynamic gradient compression for efficient communication. Prototype-based FL (e.g. FedProto [24]) collaboratively performs model training by exchanging prototypes instead of gradients.

Unfortunately, the generalization performance of these PFL methods mentioned above will decrease when data is scarce, and only a few of them [27, 28] have considered using classifier collaboration to solve this problem. Compared to them, our algorithm ensures high generalization performance in scenarios with high heterogeneous and sparse data, while also being more lightweight and efficient.

3 Notations and Preliminary

In this section, we first provide formal definitions used in PFL training in and then examine the class-imbalance problems based on our theoretical insights on simplified PFL models.

3.1 Terminology

Before illustrating our method, we firstly define the symbols that may be used later. Suppose that there are K clients and each client k holds its own dataset $(\mathbf{x}_i, y_i)_{i=1}^{N_k} \in D_k$. The global dataset $D = \cup_{k=1}^K D_k$ has C classes so that every single input is in the Cartesian product of input space \mathcal{X} and label space \mathcal{Y} .

Furthermore, we adopt the general practice in computer vision and other domains, by decomposing each client's model parameters (\mathbf{w}) into a *feature extractor* ($\boldsymbol{\theta}$) and a *classifier* (ϕ); formally, $\mathbf{w} = \{\boldsymbol{\theta}, \phi\}$ [1, 4, 2]. The classifier often indicates the last one or few fully-connected layers, while the rest layers are incorporated into feature extractor. E.g., in convolutional networks, a feature extractor typically consists of series of convolutional layers. We denote \mathbf{h} as the output of feature extractor $f_{\boldsymbol{\theta}} : \mathcal{X} \rightarrow \mathcal{H}$. We adopt another common terminology *proxy*, used in previous works [15], to indicate linear weights of different class denoted by $\phi = \{\phi_c\}_{c=1}^C$, where each proxy is denoted by ϕ_c . Besides, we use z_c to denote each *logit* for class c , which is a probability. The logits over all classes form the a vector of *logits*, denoted by $\mathbf{z} = (z_1, z_2, \dots, z_C)$, which are computed by $\mathbf{z} = f_{\phi} : \mathcal{H} \rightarrow \mathbb{R}^C$.

We consider that a total of K clients perform the PFL task, where each client k trains its own model by minimizing the local loss $\mathcal{L}_k(\cdot)$. Then the global problem can be formulated as (1):

$$\min_{\mathbf{W}} \mathcal{L}(\mathbf{W}) := \sum_k \frac{|D_k|}{|D|} \mathbb{E}_{(\mathbf{x}, y) \in D_k} [\mathcal{L}_k(\mathbf{w}_k; (\mathbf{x}, y))] \quad (1)$$

where $\mathbf{W} = \{\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_K\}$ is the collection of k model weights, different from training a single model in traditional FL. Equation (1) also naturally leads to the optimal \mathbf{w}^* for each client:

$$\mathbf{w}_k^* = \arg \min_{\mathbf{w}} \mathcal{L}_k(\mathbf{w}_k; (\mathbf{x}, y)) \quad (2)$$

For better scalability and interpretability, the local loss $\mathcal{L}_k(\cdot)$ is not explicitly specified here. It can be the cross-entropy loss function used in vanilla FL or it can incorporate unique penalty terms similar to those in knowledge distillation [24].

In addition, when describing matrices in the subsequent text, we typically use uppercase bold letters to represent them, while column vectors are represented by lowercase bold letters. The elements of matrices and vectors are denoted using regular letters.

3.2 Class-Imbalance Training

Data heterogeneity. When discussing data heterogeneity, we always refer to feature skew and label skew. Then the k -th client follows $P_k(\mathbf{x}, y) = P_k(\mathbf{x}|y)P_k(y)$ which means that the joint distribution of dataset is influenced by both conditional distributions $P_k(\mathbf{x}|y)$ and prior distributions $P_k(y)$. The variation of conditional distribution refers to the different distribution of data features under the same label, while diverse prior distributions across labels and clients reflect the class-imbalance issue in PFL. This also indirectly indicates that methods for managing feature distributions should be compatible with those dealing with label distributions.

Multi-classification problem under class-imbalance. It is important to emphasize that especially in classification problems, the classifier learns the “patterns” of features. To better illustrate the impact of class imbalance on model training, for simplicity, let's consider a linear classifier $f_{\mathbf{w}} : \mathcal{X} \rightarrow \mathbb{R}^C$. As mentioned in the previous section, \mathbf{w} is actually a matrix composed of different label-specific proxies ϕ_c . After performing vector multiplication with \mathbf{x} , it will be fed into the Softmax function to output logits:

$$\begin{aligned} \mathbf{z} &= \text{softmax}(\mathbf{w}^T \mathbf{x}) \\ &= (z_1, z_2, \dots, z_C) \\ &= \left(\frac{e^{\phi_1^T \mathbf{x}}}{\sum e^{\phi_c^T \mathbf{x}}}, \frac{e^{\phi_2^T \mathbf{x}}}{\sum e^{\phi_c^T \mathbf{x}}}, \dots, \frac{e^{\phi_C^T \mathbf{x}}}{\sum e^{\phi_c^T \mathbf{x}}} \right) \end{aligned} \quad (3)$$

When the loss function is cross-entropy, given (3), the partial derivatives with respect to each proxy ϕ_c can be derived, which will vary depending on the certain class of training samples, thus affecting the backpropagation process:

$$\phi'_c = \phi_c - \eta \frac{\partial \mathcal{L}}{\partial \phi_c} = \begin{cases} \phi_c + \eta(1 - z_{c_0})\mathbf{x}, & c_0 = c \\ \phi_c - \eta z_{c_0}\mathbf{x}, & c_0 \neq c \end{cases} \quad (4)$$

This means that when the data distribution is extremely skewed, such as in a ten-class classification problem where only data with a single label c is used for training, the update of the proxy corresponding to c follows the first case of equation (4). This leads to a significant amount of training for ϕ_c , aligning the proxy with the features of class c and gradually learning the “pattern” of class c , resulting in a tendency towards high prediction values. On the other hand, the update process for the remaining proxies follows the second case, gradually moving away from the “pattern” of class c , and failing to learn their own “patterns”.

Prediction bias. We then examine the prediction bias under class-imbalanced circumstances through the concept of prototype [24, 32]. Suppose that for any random data sample $\tilde{\mathbf{x}}_c \in \mathbb{R}^d$ belonging to a certain class $c \in [C]$ ($d > C$), there exists a prototype $\bar{\mathbf{x}}_c = \mathbb{E}_{\tilde{\mathbf{x}}_c \sim P_k(\mathbf{x}|c)} [\tilde{\mathbf{x}}_c]$. If the prototype formed by different categories are linearly independent from each other, the classifier will inevitably fail to distinguish between these two distinct concepts c_1, c_2 from random data $\tilde{\mathbf{x}}_{c_1}, \tilde{\mathbf{x}}_{c_2}$. In such cases, the dimension of the input needs to be to expand to incorporate additional information that aids in discrimination. In the other case, where the prototypes are orthogonal to each other, the rank of the matrix formed by prototypes from different categories will be equal to the number of classes C , and the prototypes will be linearly independent from each other. This is analogous to performing 0-1 classification where the data points have no correlation. In this scenario, even if only one sample is provided for each label, the classifier can perform excellently.

However, in reality, prototypes can be linearly dependent in classification tasks. If training is conducted in the presence of class-imbalance, linear classifiers continuously fit proxies to their respective prototypes, assigning higher weights when predicting their own data. Due to the presence of linear relationships among prototypes and unequal learning of different proxies by (4), it is highly likely to confuse the classes corresponding to two prototypes that have a linearly dependent relationship. This is where prediction bias arises: when clients make predictions, they often give their dominant classes larger prediction values. Conversely, *two classifiers with the same prediction bias are likely to share similar dominant classes.*

Given these preliminary understanding of potential results due to class-imbalance, we are ready to propose our algorithm in Section 4.1 and 4.2, where the above analysis will be further validated.

4 Our Method: FedReMa

Due to the class-imbalance issue in PFL training, aggregating the classifiers of clients that have very diverse expertise across classes might be detrimental. Based on the analysis in Section 3.2, a group of clients with similar biases may share a set of dominant classes in common. It is then feasible to assess their data distributions through the similarity of logits. By this means, we can efficiently identify the most relevant peers for each client and establish their training dependencies. Performing classifier collaboration based on these inter-client dependencies can potentially reduce the overfitting of proxies. Our design goal is to enhance each model's personalizability on dominant classes, preserving their preferences tailored to local datasets, while improving the generalization performance through the peers' knowledge on the client's non-dominant classes.

To realize this, we propose *FedReMa*, an adaptive method that leverages a novel relevant matching algorithm (named as MDS shown in Algorithm 1) for classifier synergy and carefully identified critical co-learning period (CCP), adapting to the dynamics of clients' prediction expertise; the feature extractor aggregation is the same as standard FedAvg. As shown in Figure 2, after each local training is completed, the server judges whether the current round is the critical co-learning period (Section 4.2) and adopts MDS if it is the case; otherwise, a distinct classifier synergy algorithm leveraging historical peer selections is performed (Section 4.3).

4.1 Relevant Matching

Local Training. As shown in lines 1-6 of Algorithm 2, like most federated learning algorithms, we assume that during communication round $t \in [1, T]$, the clients will perform multiple epochs of parameter updates. At the beginning of t , the clients receive a personalized model from the server and initialize it as $\mathbf{w}_k^{(t)} \leftarrow \mathbf{w}_k^{(t-1)}$. They then conduct E epochs of local training and transmit the updated model $\hat{\mathbf{w}}_k^{(t)} = \{\hat{\theta}_k^{(t)}, \hat{\phi}_k^{(t)}\}$ back to the server for the next step of the algorithm. Here, we have the option to employ simple stochastic gradient descent (SDG) for training, or explore different training strategies. The distinctions between various approaches will be observed in Section 5.2.

Relevant Matching. In this section, we will introduce how to perform the relevant matching algorithm using the logits similarity of the uploaded classifier set $\{\phi_k^{(t)}\}_{k \in \mathcal{K}}$. For the sake of brevity, we will ignore the superscript (t) of some variables.

Based on the analysis from Section 3.2, due to training under class-imbalance, there is an inequality in the sufficiency of training for proxies corresponding to different classes. This leads to the models in

Algorithm 1: Maximum difference segmentation (MDS)

Input: Relativity vector $\mathbf{s} = (s_1, \dots, s_K)$

Output: Set of the most relevant nodes \mathcal{A} and maximum difference Δs^*

- 1 Sort the relativity element:
 $\text{idx} = \text{SortIndices}(\mathbf{s}) = (\text{idx}_1, \dots, \text{idx}_K)$,
 $\text{Sort}(\mathbf{s}) = (s_{\text{idx}_1}, \dots, s_{\text{idx}_K})$;
 - 2 Calculate difference $\Delta s_{\text{idx}_i} = s_{\text{idx}_{i+1}} - s_{\text{idx}_i}$;
 - 3 Get $m = \arg \max_{i \in [K-1]} \Delta s_{\text{idx}_i}$ $\Delta s^* = \Delta s_{\text{idx}_m}$;
 - 4 **Return** $\mathcal{A} = \{\text{idx}_m, \text{idx}_{m+1}, \dots, \text{idx}_K\}$ and Δs^* ;
-

the early stages being prone to prediction bias. Even when we input a randomly generated feature, different classifiers will provide biased logits that are strongly correlated with the local class distribution.

We test each client's classifier response by inputting a feature $\tilde{\mathbf{h}}$ where each element is independently sampled from a uniform distribution in the range $[0, 1)$. We then construct a soft prediction method, using $\mathbf{p}_k = \text{softmax}(\phi_k^\top \tilde{\mathbf{h}}/M)$ to control the sharpness of logits via a temperature parameter M (Unless otherwise specified, we set $M = 0.5$). We establish this by borrowing the idea from the concept of response-based knowledge in knowledge distillation [7].

The new \mathbf{p}_k we call it *soft logits*. Consider a connected graph \mathcal{S} that can be fully represented by a relativity matrix $\mathbf{S} \in \mathbb{R}^{K \times K}$, where the weights are obtained through the computation of cosine similarity between the soft logits of different clients:

$$s_{k_1, k_2} = \frac{\mathbf{p}_{k_1} \cdot \mathbf{p}_{k_2}}{\|\mathbf{p}_{k_1}\| \|\mathbf{p}_{k_2}\|} \quad (5)$$

Due to the properties of the softmax function, the similarity values calculated in the end are in the range of $[0, 1]$. For each client k , there exists a relevance vector $\mathbf{s}_k = (s_{k,1}, \dots, s_{k,K})$. We employ the *maximum difference segmentation* method (MDS, shown as Algorithm 1) to obtain the set of most relevant clients \mathcal{A}_k for client k , using the maximum difference as the boundary for client selection.

Then, the server performs distinct aggregation methods on the uploaded feature extractors and classifiers:

$$\theta^{(t)} \leftarrow \frac{1}{|D|} \sum_{i \in \mathcal{K}} |D_i| \hat{\theta}_i^{(t)} \quad (6)$$

$$\phi_k^{(t)} \leftarrow \frac{1}{|D_{\mathcal{A}_k^{(t)}}|} \sum_{i \in \mathcal{A}_k^{(t)}} |D_i| \hat{\phi}_i^{(t)} \quad (7)$$

The server aggregates a robust global feature extractor in equation (6), while equation (7) ensures the personalized aggregation of classifiers.

4.2 Critical Co-learning Period

As shown in Figure 3, we gain another key finding through experiments: *As the training progresses, the effectiveness of relevant matching degrades.* This can imply that models in the later stages of training already have the ability to identify non-dominant classes. The model performance might be undermined we continuously adopt the previous metric and method for classifier synergy. Therefore, we propose a critical co-learning period (CCP) method based on the historical maximum similarity difference given by MDS. We define the average maximum similarity difference as the average value of $\{\Delta s_k^t\}_{k \in \mathcal{K}}$ for all clients in round t .

$$\Delta \bar{s}^t = \sum_{k \in \mathcal{K}} \Delta s_k^t \quad (8)$$

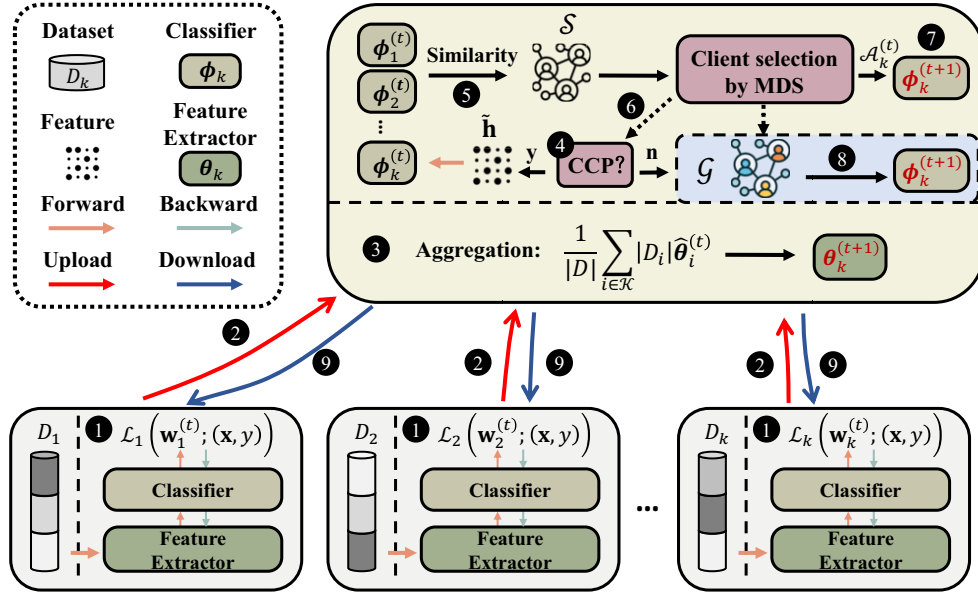


Figure 2: Workflow of FedReMa. (1) Local training. (2) Upload model to server. (3) Aggregate global feature extractors. (4) Determine whether the current period is a critical learning period, if so, go to (5), otherwise go to (8). (5) Measure the task relevance \mathcal{S} using similarity of their logits, then perform MDS to filter out irrelevant clients. At the same time, the server will record the client's historical dependencies \mathcal{G} . (6) Update CCP status. (7) Aggregate personalized classifiers by $\mathcal{A}_k^{(t)}$. (8) Aggregate personalized classifiers based on historical dependencies \mathcal{G} . (9) Server communicates personalized models to all clients.

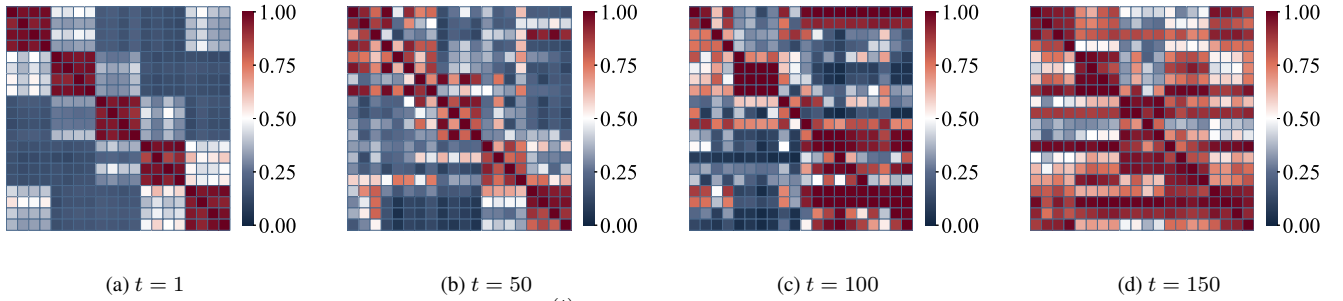


Figure 3: The heatmap of $\mathbf{S}^{(t)}$ on different communication round t (FedAvg)

The critical co-learning period can be determined via (9):

$$CCP = \begin{cases} True, & \frac{\Delta \bar{s}^t}{\max_{\tau \in [t]} \Delta \bar{s}^\tau} > \delta \\ False, & \text{otherwise} \end{cases} \quad (9)$$

Here, δ is a constant threshold. It is evident that the determination of CCP is triggered by the decrease of $\Delta \bar{s}^t$. Once we are unable to differentiate relevant clients based on similarities, the co-learning in this stage becomes ineffective. If the MDS algorithm continues to be applied, there will be a degradation in accuracy. Another way to effectively aggregate the classifiers must be used.

4.3 Personalization via Historical Peer Matching

The primary purpose of setting the critical co-learning period is (1) to simplify the computational cost in the later phase, and (2) to maintain personalization. During the CCP , server maintains a dependency map $\mathcal{G} \in \mathbb{R}^{K \times K}$, where each column vector $\mathbf{g}_k = (g_{k,1}, \dots, g_{k,K})^\top$ represents the number of times client k selects other clients as its relevant peers. When client j is repeatedly selected

by client i , a high value of $g_{i,j}$ will indicate a strong dependency relationship of i on j .

When CCP is determined to be *False*, we adopt an algorithm that leverages the historical matching decisions. The rationale is that: We can only utilize the MDS algorithm for client selection to facilitate classifier collaboration during CCP . However, after CCP , the clients may have accumulated a certain level of knowledge, causing the distinction between soft logits of relevant clients to become ambiguous. Thus, we \mathcal{G} is a historical dependency map that utilizes the $\mathcal{A}_k^{(t)}$ when the most relevant clients is still distinguishable. It records the number of times a particular client selects others, which itself serves as a form of weight information.

$$\phi_k^{(t)} \leftarrow \frac{1}{\sum_{i \in \mathcal{K}} g_{k,i}} \sum_{i \in \mathcal{K}} g_{k,i} \cdot \hat{\phi}_i^{(t)} \quad (10)$$

With the help of dependency map \mathcal{G} , we can perform classifier collaboration directly via Equation (10). In the experiments presented later in the Section 5.2, we can observe that this design effectively reduces computational complexity and training time.

Algorithm 2: PFL via the Most Relevant Peers (*FedReMa*)

Input: initial model $\mathbf{w} = \{\theta, \phi\}$, client number K , global round T , local epoch E

Output: Personalized model \mathbf{w}_k^*

```

1 for  $t = 1$  to  $T$  do
2   foreach client  $k \in \mathcal{K}$  in parallel do
3     Initialize the local model  $\mathbf{w}_k^{(t)} \leftarrow \mathbf{w}_k^{(t-1)}$ ;
4      $\hat{\mathbf{w}}_k^{(t)} \leftarrow \text{ClientTrain}(\mathbf{w}_k; D_k)$ ;
5     Communicate  $\hat{\mathbf{w}}_k^{(t)}$  back to the server;
6   end
7   Server calculates  $\mathbf{S}^{(t)}$ ;
8   Server aggregates the feature extractor by Eq. (6);
9   if CCP then
10    for each client  $k \in \mathcal{K}$  do
11      Get relevant set  $\mathcal{A}_k^{(t)} \leftarrow \text{MDS}(\mathbf{s}_k^{(t)})$ ;
12      Personalization by Eq. (7);
13    end
14    Server counts  $\mathcal{G}$ ;
15    Check CCP by Eq. (9);
16  else
17    Personalization by Eq. (10);
18  end
19  Server communicates  $\mathbf{w}_k^{(t)} = \{\theta^{(t)}, \phi_k^{(t)}\}$  to all clients;
20 end

```

5 Experiment

In this section we will present our experimental results.

5.1 Settings

Dataset. We evaluate the algorithm using three classic computer vision classification datasets: MNIST, FMNIST, and CIFAR-10. These datasets consist of 50000 training images (60000 For FMNIST) with 10 classes. To control data heterogeneity across local datasets, we adopt the same data partitioning scheme as [10, 28]. Specifically, while ensuring dataset equality, we divide all the clients into five groups. In each group, three labels are designated as dominant labels, with a proportion parameter $s\%$ controlling the fraction of IID data, while the remaining $(1 - s)\%$ data belong to dominant labels.

Baselines. We compare the following methods with our *FedReMa*: (1) Local (no model aggregation), (2) FedAvg [18], (3) FedProx [13], (4) MOON [12], (5) FedProto [24], (6) FedPer [1], (7) PerFedAvg [5], (8) FedBABU [20], (9) FedPAC [28], (10) FedDistill [9].

Model Settings. For MNIST and FMNIST, we use a CNN model consisting of 2 convolutional layers and 2 fully connected layers. For CIFAR-10, we use a simplified AlexNet for training. Specifically, we have reduced the size of the convolutional kernels and omitted the use of Dropout and Local Response Normalization (LRN) layers. This version of AlexNet is better suited for CIFAR-10. We consider the fully connected layers of each network as classifiers, and the remaining convolutional layers as feature extractors.

Hyperparameters. We use mini-batch stochastic gradient descent (SGD) as the optimizer for all methods. The number of communication rounds is set to $T = 500$, and the number of local training epochs between consecutive communications are set to $E = 5$. The batch size is set to 100, and the learning rate is set to $\eta = 0.01$. For FedProx, μ is set to 1. For MOON, μ is set to 1 (5 for the CIFAR-10). The regularization λ is 1 for both FedProto and FedPAC. For other

hyperparameters not mentioned, we adopt the best of them reported in their respective original papers.

Evaluation Metrics. In contrast to conventional FL where the global is finally tested on a centralized dataset, PFL measures the personalized accuracy over clients under individual local datasets. All experiments were conducted on a lab server with a NVIDIA GeForce RTX 3090 GPU.

5.2 Performance Evaluation

Comparison across different datasets. To simulate a highly heterogeneous environment, we initially set the sample size for each client to be 600 and divided it via a 4:1 ratio for both training and testing. We conducted experiments with $s = 0.2$. Each set of experiments were repeated 5 times with different random seeds. Table 1 shows the performance comparison. Our *FedReMa* achieves the highest accuracy in the majority of settings. It is worth noting that FedAvg remains a strong baseline, exhibiting good performance even in scenarios with high heterogeneity and limited samples.

Table 1: Accuracy on different datasets.

	#Method	MNIST	FMNIST	CIFAR-10
TFL	Local	94.3±0.2%	86.0±0.2%	58.9±0.1%
	FedAvg	98.4±0.1%	86.1±0.3%	61.9±0.6%
	FedProx	97.8±0.2%	84.8±0.1%	61.6±0.8%
	MOON	98.4±0.2%	86.0±0.2%	61.9±0.7%
PFL	FedProto	95.6±0.1%	80.8±0.4%	56.1±0.4%
	FedPer	94.8±0.1%	86.1±0.1%	59.5±0.6%
	Per-FedAvg	97.6±0.3%	88.1±0.3%	62.3±2.1%
	FedDistill	96.3±0.1%	87.0±0.1%	61.2±0.5%
	FedBABU	97.1±0.2%	83.1±0.2%	53.6±0.8%
	FedPAC	98.4±0.1%	86.9±0.4%	64.9±0.1%
Ours	FedReMa	98.5±0.1%	88.2±0.3%	65.4±0.3%

Comparison of varying data heterogeneities. We further conducted experiments varying data heterogeneity, shown in Table 2. Each client had only 100 samples, imposing limitations on model generalization.

FedReMa still maintains the best performance when different values of s are chosen. Note that the generalization performance of the algorithms is greatly challenged in the case of sparse data. When the data heterogeneity is high ($s = 0.2$), various methods can still maintain a certain level of accuracy. FedPer and FedPAC perform well, while FedProto’s accuracy is significantly affected, even experiencing a complete collapse in the later stages of training (Figure 4a, 4b and 4c). We believe this is due to the inability to achieve unbiased estimation of prototypes with sparse data. As the data homogeneous increases, FL methods stand out while PFL methods experience varying degrees of accuracy degradation.

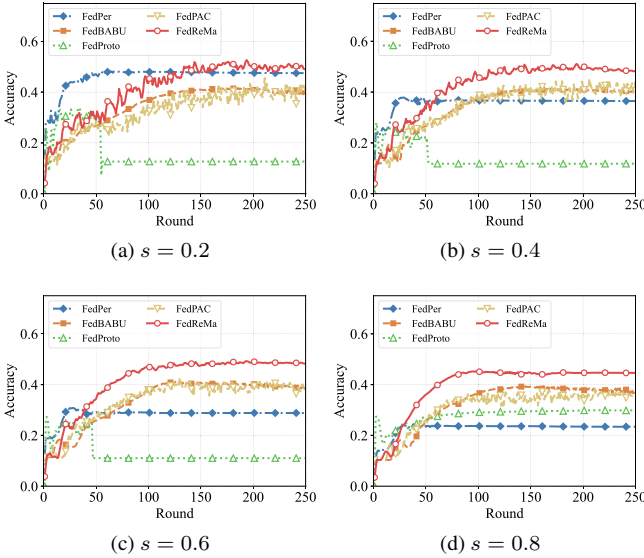
However, when the data approaches homogeneous, *FedReMa* also experiences performance degradation. It is reasonable because we do not introduce additional techniques apart from adaptive personalized aggregation of classifiers, for reducing time complexity. Therefore, under less heterogeneous data, the CCP ends prematurely, resulting in insufficient personalization.

Sensitivity. We also conducted experiments varying δ , our algorithm’s key hyperparameter, which controls the epochs of the CCP. The experimental results are shown in the Table 3. “Final accuracy” refers to the average accuracy of the last 5 communication rounds.

The best accuracy of *FedReMa* doesn’t vary significantly when δ increases, as the best model occurs during the early phase when the MDS algorithm is still being executed. When $\delta = 0.1$, the running

Table 2: Accuracy on CIFAR-10 with different data heterogeneity.

#Method	s=0.2	s=0.4	s=0.6	s=0.8
Local	48.1±0.8%	38.1±0.6%	30.2±0.4%	24.2±0.5%
FedAvg	46.4±0.9%	47.1±1.2%	47.4±1.2%	45.4±1.3%
FedProx	40.7±0.6%	41.9±0.5%	41.6±0.8%	41.3±0.4%
MOON	46.4±0.8%	47.1±1.2%	47.4±1.1%	45.5±1.2%
FedProto	36.7±3.0%	29.2±0.8%	30.9±2.7%	30.6±0.9%
FedPer	47.6±0.8%	37.8±0.9%	30.9±0.9%	25.1±0.6%
FedBABU	41.3±0.5%	41.1±0.5%	40.9±0.6%	40.1±0.7%
FedPAC	48.0±2.1%	45.7±1.2%	41.6±0.8%	38.3±0.3%
FedReMa	52.9±1.2%	50.2±0.7%	47.6±1.4%	46.5±1.8%

**Figure 4:** The test accuracy while training process.

time for *FedReMa* is 1403.08s. While for $\delta = 0.9$, it is 760.89s (54.23%). Maintaining personalization can be observed from the final accuracy decrease. An unreasonably small δ can lead to degradation to FedAvg, while a too large δ may fail to provide a good enough \mathcal{G} in the second stage. Overall, although our algorithm is relatively robust to the value of δ , $\delta = 0.5$ is the best choice in this experimental configuration.

Compatibility with FedPAC. The previous experiments demonstrated the excellent performance of *FedReMa* with the help of MDS algorithm and critical co-learning period. However, *FedReMa* itself is a server-side algorithm and does not introduce any new methods on the client side. Thus we have conducted additional testing on the FedPAC method, which has also undergone classifier collaboration. We can confirm that if we combine the techniques in FedPAC with our *FedReMa*, the performance can be further enhanced.

The training strategies on client-side in FedPAC are as follows:

1. **Feature alignment:** As shown in (11), an additional feature alignment penalty term is added in FedPAC.

$$\begin{aligned} \mathbf{w}_k^* &= \arg \min_{\mathbf{w}} \mathcal{L}_k(\mathbf{w}_k; (\mathbf{x}, y)) \\ &= \arg \min_{\mathbf{w}} \mathcal{L}_{CE}(\mathbf{w}_k; (\mathbf{x}, y)) + \lambda \mathcal{L}_d(\bar{h}; h) \end{aligned} \quad (11)$$

Table 3: Sensitivity of δ .

δ	0.1	0.3	0.5	0.7	0.9
Best Acc	63.36%	64.13%	65.40%	65.26%	65.06%
Final Acc	61.64%	60.73%	63.77%	62.30%	63.03%
Epochs of CCP	500	132	36	16	7

Table 4: Accuracy and Runtime Comparison with FedPAC and extended versions of our *FedReMa*.

#Method	Accuracy	Runtime
FedReMa	65.40±0.32%	796.5s
FedPAC	64.93±0.13%	4464.3s
FedReMa+F	66.31±0.55%	>3000s
FedReMa+A&F	69.73±0.98%	>3000s

where \mathcal{L}_d is usually the L2 norm. This approach requires additional communication and feature aggregation on the server side.

2. **Alternately train feature extractors and classifiers:** One advantage of alternating optimization (optimizing the classifier first, then the feature extractor) is that each optimization step of the classifier is performed in the same feature space. Thus, comparing the logits of the classifier becomes more meaningful.

It is not difficult to see that these are all strategies that are beneficial for normalizing features. By incorporating these client training strategies, we aim to enhance the performance of *FedReMa* and showcase the potential for even better results. The result is shown in Table 4, where "F" means *Feature Alignment* and "A" means *Alternately Training*. Those strategies indeed lead to better performance. Compared with FedPAC, our method has a huge advantage in computational complexity.

Ablation study. We finally verify the effectiveness of each component of our algorithm. As shown in the Table 5, the two components we designed complement each other and both achieve satisfactory accuracy. This means that both components are well suited to the task of parameter personalization.

Table 5: Ablation experiment.

#Method	Best Acc	Final Acc	Runtime
None(FedAvg)	61.95±0.65%	61.80%	691.58s
+MDS	63.36±0.79%	61.64%	1403.08s
+MDS&CCP	65.40±0.32%	63.77%	796.54s

6 Conclusion

In this paper, we proposed *FedReMa*, an improved personalized federated learning framework by strategically leveraging the expertise of different clients on specific data classes throughout the training process. *FedReMa* addresses the issue of class-imbalance by dynamically identifying and harnessing clients' expertise in different phases. By performing a novel method that combines relevant matching among clients using similarities of clients' logits and the judgement of critical co-learning period, the server can efficiently identify clients with similar tasks and enhance personalized accuracy through the personalized aggregation of classifiers. *FedReMa* outperforms existing methods significantly across various personalized federated learning scenarios.

There are several possible future research directions. For example, extending the algorithm to more general tasks, providing stronger theoretical analysis for the algorithm, etc. Another example is to explore how the correlation between features of different categories will affect the update of the classifier, and maybe there is a possible way to optimize the aggregation of the classifier by linking the correlation between features. Furthermore, there are many interesting potential optimization directions to extend our algorithm. For example, we can employ more fine-grained aggregation of different proxies of the classifier and layer-by-layer personalization based on feature similarities across clients.

Acknowledgements

This work was supported in part by Guangdong Basic and Applied Basic Research Foundation (2024A1515010161, 2023A1515012982, 2023B1515120058, 2021B151520008), Guangzhou Basic and Applied Basic Research Program (2024A04J6367), NSFC grant (62102460), Young Outstanding Award under the Zhujiang Talent Plan of Guangdong Province, and the Ministry of Education, Singapore, under its Academic Research Fund (AcRF RG91/22) and NTU Startup.

All experiments were conducted at RTAI cluster in School of Computer Science and Engineering, Sun Yat-sen University.

References

- [1] M. G. Arivazhagan, V. Aggarwal, A. K. Singh, and S. Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019.
- [2] H.-Y. Chen and W.-L. Chao. On bridging generic and personalized federated learning for image classification. 2022.
- [3] Y. J. Cho, J. Wang, and G. Joshi. Client selection in federated learning: Convergence analysis and power-of-choice selection strategies. *arXiv preprint arXiv:2010.01243*, 2020.
- [4] L. Collins, H. Hassani, A. Mokhtari, and S. Shakkottai. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, pages 2089–2099. PMLR, 2021.
- [5] A. Fallah, A. Mokhtari, and A. Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems*, 33:3557–3568, 2020.
- [6] J. Goetz, K. Malik, D. Bui, S. Moon, H. Liu, and A. Kumar. Active federated learning. *arXiv preprint arXiv:1909.12641*, 2019.
- [7] J. Gou, B. Yu, S. J. Maybank, and D. Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021.
- [8] J. Jang, H. Ha, D. Jung, and S. Yoon. Fedclassavg: Local representation learning for personalized federated learning on heterogeneous neural networks. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–10, 2022.
- [9] E. Jeong, S. Oh, H. Kim, J. Park, M. Bennis, and S.-L. Kim. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. *arXiv preprint arXiv:1811.11479*, 2018.
- [10] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh. Scaffold: Stochastic controlled cveraging for federated learning. In *International conference on machine learning*, pages 5132–5143. PMLR, 2020.
- [11] B. Kim, H. Kim, K. Kim, S. Kim, and J. Kim. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9012–9020, 2019.
- [12] Q. Li, B. He, and D. Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10713–10722, 2021.
- [13] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.
- [14] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang. On the convergence of fedavg on non-iid data. In *Proc. of International Conference on Learning Representations*, 2019.
- [15] X.-C. Li and D.-C. 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*, pages 995–1005, 2021.
- [16] P. P. Liang, T. Liu, L. Ziyin, N. B. Allen, R. P. Auerbach, D. Brent, R. Salakhutdinov, and L.-P. Morency. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*, 2020.
- [17] Y. Liu, Y. Shi, Q. Li, B. Wu, X. Wang, and L. Shen. Decentralized directed collaboration for personalized federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23168–23178, 2024.
- [18] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [19] Y. Mei, B. Guo, D. Xiao, and W. Wu. Fedvfv: Personalized federated learning based on layer-wise parameter updates with variable frequency. In *2021 IEEE International Performance, Computing, and Communications Conference (IPCCC)*, pages 1–9. IEEE, 2021.
- [20] J. H. OH, S. Kim, and S. Yun. Fedbabu: Toward enhanced representation for federated image classification. In *10th International Conference on Learning Representations, ICLR 2022. International Conference on Learning Representations (ICLR)*, 2022.
- [21] X. Shang, Y. Lu, G. Huang, and H. Wang. Federated learning on heterogeneous and long-tailed data via classifier re-training with federated features. *arXiv preprint arXiv:2204.13399*, 2022.
- [22] C. T. Dinh, N. Tran, and J. Nguyen. Personalized federated learning with moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020.
- [23] R. Tamirisa, C. Xie, W. Bao, A. Zhou, R. Arel, and A. Shamsian. Fed-select: Personalized federated learning with customized selection of parameters for fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23985–23994, 2024.
- [24] Y. Tan, G. Long, L. Liu, T. Zhou, Q. Lu, J. Jiang, and C. Zhang. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 8432–8440, 2022.
- [25] J. Wang and M.-L. Zhang. Towards mitigating the class-imbalance problem for partial label learning. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2427–2436, 2018.
- [26] C. Wu, F. Wu, L. Lyu, Y. Huang, and X. Xie. Communication-efficient federated learning via knowledge distillation. *Nature communications*, 13(1):2032, 2022.
- [27] J. Xu, Y. Yan, and S.-L. Huang. Fedper++: Toward improved personalized federated learning on heterogeneous and imbalanced data. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 01–08. IEEE, 2022.
- [28] J. Xu, X. Tong, and S.-L. Huang. Personalized federated learning with feature alignment and classifier collaboration. *arXiv preprint arXiv:2306.11867*, 2023.
- [29] Y. Yang and Z. Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in neural information processing systems*, 33:19290–19301, 2020.
- [30] Y. Zhang, T. Gao, C. Li, and C. W. Tan. Coded federated learning for communication-efficient edge computing: A survey. *IEEE Open Journal of the Communications Society*, 2024.
- [31] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.
- [32] T. Zhou, J. Zhang, and D. H. Tsang. Fedfa: Federated learning with feature anchors to align features and classifiers for heterogeneous data. *IEEE Transactions on Mobile Computing*, 2023.
- [33] H. Zhu, J. Xu, S. Liu, and Y. Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021.