

FedCP: Separating Feature Information for Personalized Federated Learning via Conditional Policy

Jianqing Zhang
Shanghai Jiao Tong University
Shanghai, China
tsingz@sjtu.edu.cn

Yang Hua
Queen's University Belfast
Belfast, UK
y.hua@qub.ac.uk

Hao Wang
Louisiana State University
Baton Rouge, USA
haowang@lsu.edu

Tao Song
Shanghai Jiao Tong University
Shanghai, China
songt333@sjtu.edu.cn

Zhengui Xue
Shanghai Jiao Tong University
Shanghai, China
zhenguixue@sjtu.edu.cn

Ruhui Ma
Shanghai Jiao Tong University
Shanghai, China
ruhuima@sjtu.edu.cn

Haibing Guan
Shanghai Jiao Tong University
Shanghai, China
hbguan@sjtu.edu.cn

ABSTRACT

Recently, personalized federated learning (pFL) has attracted increasing attention in privacy protection, collaborative learning, and tackling statistical heterogeneity among clients, *e.g.*, hospitals, mobile smartphones, *etc.* Most existing pFL methods focus on exploiting the global information and personalized information in the client-level model parameters while neglecting that data is the source of these two kinds of information. To address this, we propose the **Federated Conditional Policy (FedCP)** method, which generates a conditional policy for each sample to separate the global information and personalized information in its features and then processes them by a global head and a personalized head, respectively. FedCP is more fine-grained to consider personalization in a sample-specific manner than existing pFL methods. Extensive experiments in computer vision and natural language processing domains show that FedCP outperforms eleven state-of-the-art methods by up to 6.69%. Furthermore, FedCP maintains its superiority when some clients accidentally drop out, which frequently happens in mobile settings. Our code is public at <https://github.com/TsingZ0/FedCP>.

CCS CONCEPTS

• **Computing methodologies** → **Multi-agent systems; Distributed algorithms; Supervised learning.**

KEYWORDS

Federated Learning; Statistical Heterogeneity; Personalization; Conditional Computing; Feature Separation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0103-0/23/08...\$15.00

<https://doi.org/10.1145/3580305.3599345>

ACM Reference Format:

Jianqing Zhang, Yang Hua, Hao Wang, Tao Song, Zhengui Xue, Ruhui Ma, and Haibing Guan. 2023. FedCP: Separating Feature Information for Personalized Federated Learning via Conditional Policy. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3580305.3599345>

1 INTRODUCTION

Nowadays, many web-based services, such as personalized recommendations [53, 55, 57], benefit from artificial intelligence (AI) and the huge volume of data generated locally on various clients [19], *e.g.*, hospitals, mobile smartphones, internet of things, *etc.* At the same time, legislation endeavors on data privacy protection continue to increase, *e.g.*, General Data Protection Regulation (GDPR) of Europe [38] and California Consumer Privacy Act (CCPA) [7]. Due to privacy concerns and regulations, centralized AI faces significant challenges [34, 49]. On the other hand, because of the data sparsity problem, it is hard to learn a reasonable model for a given task independently on each client [19, 25, 43].

Federated learning (FL) is proposed as a collaborative learning paradigm [19, 32, 51, 54] to utilize local data on the participating clients for the global model training without sharing the private data of clients. As one of the famous FL methods, FedAvg conducts four steps in each communication iteration: (1) The server sends the old global model parameters to the selected clients. (2) Each selected client initializes the local model with the received global parameters and trains the local model on local data. (3) The selected clients upload the updated local model parameters to the server. (4) The server generates new global model parameters by aggregating the received client model parameters. However, in practice, the data on the client is typically not independent and identically distributed (non-IID) as well as unbalanced [19, 25, 50, 54]. With this statistical heterogeneity challenge [25, 43], the single global model in traditional FL methods, such as FedAvg, can hardly fit the local data well on each client and achieve good performance [15, 42].

To meet the personalized demand of each client and address the challenge of statistical heterogeneity in FL, personalized federated learning (pFL) comes along that focuses on learning personalized models rather than a single global model [27, 42]. Most existing pFL methods consider the global model as a container that stores the global information and enriches the personalized models with the parameters in the global model. However, they only focus on client-level model parameters, *i.e.*, the global/personalized model to exploit the global/personalized information. Specifically, the meta-learning-based methods (such as Per-FedAvg [8]) only fine-tune global model parameters to fit local data, and the regularization-based methods (such as pFedMe [42], FedAMP [15], and Ditto [24]) only regularize model parameters during local training. Although personalized-head-based methods (such as FedPer[2], FedRep [6], and FedRoD [4]) explicitly split a backbone into a global part (feature extractor) and a personalized part (head), they still focus on exploiting global and personalized information in model parameters rather than the source of information: **data**. As the model is trained on data, **the global/personalized information in model parameters is derived from client data**. In other words, the heterogeneous data on clients contains both global and personalized information. As shown in Figure 1, widely-used colors, *e.g.*, blue, and rarely-used colors, *e.g.*, purple and pink, contain global information and personalized information in images, respectively.

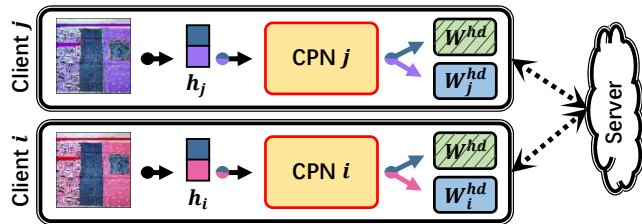


Figure 1: An example for FedCP. $h_{i/j}$: extracted feature vector, CPN i/j : Conditional Policy Network, W^{hd} : frozen global head, $W_{i/j}^{hd}$: personalized head. Best viewed in color.

To exploit the global and personalized information in the data separately, we propose a **Federated Conditional Policy (FedCP)** method based on conditional computing techniques [11, 35]. Since the dimension of raw input data is much larger than the feature vector extracted by the feature extractor, we focus on the feature vector for efficiency. As the proportion of the global and personalized information in the features differ among samples and clients, we propose an auxiliary **Conditional Policy Network (CPN)** to generate the sample-specific policy for feature information separation. Then, we process the global feature information and personalized feature information by a global head and a personalized head in different routes, respectively, as shown in Figure 1. We store the personalized information in the personalized head and reserve the global information by freezing the global head without locally training it. Through end-to-end learning, CPN automatically learns to generate the sample-specific policy. We visualize six cases in Section 5.1 to show the effectiveness of the feature information separation ability.

To evaluate FedCP, we conduct extensive experiments on various datasets in two widely-used scenarios [23, 32], *i.e.*, the pathological

settings and the practical settings. FedCP outperforms eleven state-of-the-art (SOTA) methods in both scenarios, and we analyze the reasons in Section 6.1. In summary, our key contributions are:

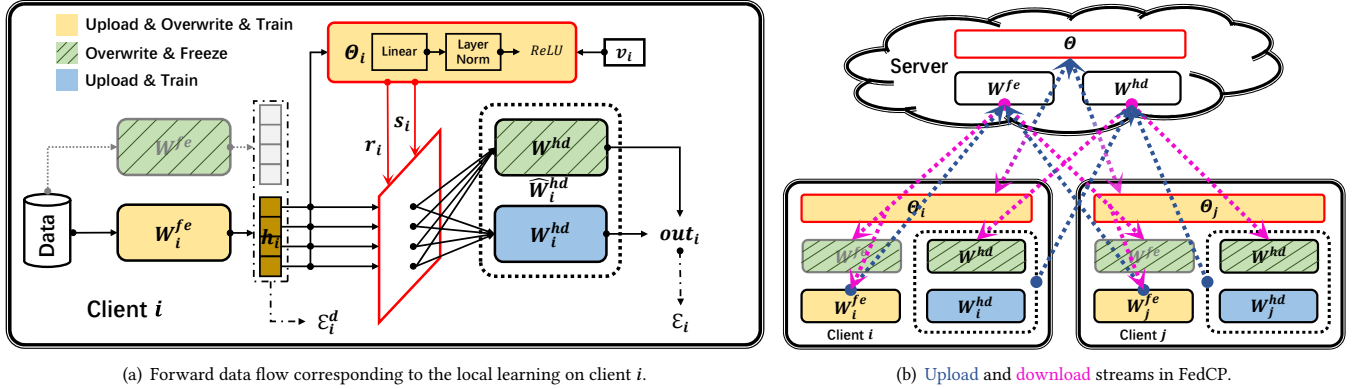
- To the best of our knowledge, we are the first to consider personalization on the sample-specific feature information in FL. It is more fine-grained than using the client-level model parameters in most existing FL methods.
- We propose a novel FedCP that generates a sample-specific policy to separate the global information and personalized information in features on each client. It processes these two kinds of feature information through a frozen global head and a personalized head on each client, respectively.
- We conduct extensive experiments in computer vision (CV) and natural language processing (NLP) domains to show the effectiveness of FedCP. Besides, FedCP keeps its superior performance even when some clients accidentally drop out.

2 RELATED WORK

2.1 Personalized Federated Learning

To collaboratively learn models among clients on their local private data while protecting privacy, traditional FL methods, such as FedAvg [32] and FedProx [26], come along. Based on FedAvg, FedProx improves the stability of the FL process through a regularization term. However, in practice, statistical heterogeneity widely exists in the FL setting, so it is hard to learn a single global model that fits well with the local data in each client [15, 19, 42].

Recently, pFL has attracted increasing attention for its ability to tackle statistical heterogeneity in FL [12, 19]. Among **meta-learning-based methods**, Per-FedAvg [8] learns an initial shared model as the global model that satisfies the learning trend for each client. Among **regularization-based methods**, pFedMe [42] learns an additional personalized model locally for each client with Moreau envelopes. In addition to learning only one global model for all clients, FedAMP [15] generates one server model for one client through the attention-inducing function to find similar clients. In Ditto [24], each client learns its personalized model locally with a proximal term to fetch global information from global model parameters. Among **personalized-head-based methods**, FedPer[2] and FedRep [6] learn a global feature extractor and a client-specific head. The former locally trains the head with the feature extractor, while the latter locally fine-tunes the head until convergence before training the feature extractor in each iteration. To bridge traditional FL and pFL, FedRoD [4] explicitly learns two prediction tasks with a global feature extractor and two heads. It uses the balanced softmax (BSM) loss [39] for the global prediction task and processes the personalized task by the personalized head. Among **other pFL methods**, FedFomo [56] calculates the client-specific weights for aggregation on each client using the personalized models from other clients. FedPHP [27] locally aggregates the global model and the old personalized model using a moving average to keep the historical personalized information. It also transfers the information in the global feature extractor through the widely-used maximum mean discrepancy (MMD) loss [10, 37]. These above pFL methods only focus on exploiting global and personalized information of model parameters but do not dig deep into data.

(a) Forward data flow corresponding to the local learning on client i .

(b) Upload and download streams in FedCP.

Figure 2: (a) The conditional policy separates information from h_i into $r_i \odot h_i$ and $s_i \odot h_i$ in the red rhomboid. Except for the feature vectors and vector v_i , a standard rectangle and a rounded rectangle represent a layer and a module, respectively. The rounded rectangle with the dashed border is \widehat{W}_i^{hd} in Eq. (6). W^{fe} (gray border) is not a part of the personalized model, where data only flows forward during training. Data flows in all the lines during training, but it only flows in the solid lines during inference. (b) For clarity, we separately show the upload and download streams for the feature extractors, the heads, and the CPNs. Still, we upload or download them as a union between the server and each client in practice. Best viewed in color.

2.2 Conditional Computing

Conditional computing is a technique that introduces dynamic characteristics into models according to task-dependent conditional inputs [11, 30, 35]. Formally, given a conditional input C (e.g., image/text, model parameter vector, or other auxiliary information) and an auxiliary module $AM(\cdot; \theta)$, a signal S can be generated by $S = AM(C; \theta)$ and used to interfere with models, such as dynamic routing and feature adaptation.

To activate specific parts in a model and process the data in different routes for each input sample, many approaches generate sample-specific policies for route selection. Conditioned on the input image, ConvNet-AIG [44] can decide which layers are needed during inference using Gumbel Softmax [17]. With a policy network, SpotTune [11] makes decisions for each image to select which blocks in a pre-trained residual network to fine-tune.

Instead of focusing on dynamic model topology, some methods propose adapting the learned features. In the few-shot learning field, TADAM [35] adapts the features through an affine transformation conditioned by the extracted task representation. In the video object detection field, TMA [45] proposes a learnable affine transformation conditioned by video frames for feature adaptation.

The above methods use conditional computing techniques but are designed for centralized AI scenarios and specific tasks. Combining the ideas of dynamic routing and feature adaptation, we devise the CPN module in our FedCP to separate global feature information and personalized feature information then process them in different routes for pFL scenarios and various tasks.

3 METHOD

3.1 Overview

In statistically heterogeneous pFL settings, non-IID and unbalanced data exist on N clients, who train their personalized models W_1, \dots, W_N in a collaborative manner. N clients own private

datasets $\mathcal{D}_1, \dots, \mathcal{D}_N$, respectively, which are sampled from N distinct distributions without overlapping.

Similar to FedPer [2], FedRep [6], and FedRoD [4], we split the backbone into a feature extractor $f: \mathbb{R}^D \rightarrow \mathbb{R}^K$, that maps input samples to feature space and a head $g: \mathbb{R}^K \rightarrow \mathbb{R}^C$, which maps from low-dimensional feature space to a label space. Following FedRep, we consider the last fully connected (FC) layer in each given backbone as the head. D , K , and C are the dimension of the input space, feature space, and label space, respectively. K is determined by the given backbone and typically $D \gg K$.

Different from FedPer, FedRep and FedRoD, on client i , we have a global feature extractor (parameterized by W^{fe}), a global head (parameterized by W^{hd}), a personalized feature extractor (parameterized by W_i^{fe}), a personalized head (parameterized by W_i^{hd}), and a CPN (parameterized by Θ_i). Specifically, **for the feature extractors**, we initialize W_i^{fe} by overwriting it with corresponding global parameters W^{fe} in each iteration, and then locally learn the personalized feature extractor. The feature generated by the changing personalized feature extractor may not fit the frozen global head during local learning. Thus, we freeze the global feature extractor after receiving and align the features outputted by the personalized feature extractor to the ones generated by the global feature extractor through the MMD loss, as shown in Figure 2(a). **For the global head**, we freeze it after it has been initialized by W^{hd} to preserve global information. In short, at the start of each iteration, we overwrite W_i^{fe} by new W^{fe} then freeze W^{fe} and W^{hd} . As shown by the non-transparent module in Figure 2(a), the personalized model used for inference (parameterized by W_i) consists of the personalized feature extractor, the global head, the personalized head, and the CPN, i.e., $W_i := \{W_i^{fe}, W^{hd}, W_i^{hd}, \Theta_i\}$. The frozen global feature extractor is only used for local learning and is not part of the personalized model. We omit iteration notation, sample index notation, and biases for simplicity. Given the local loss \mathcal{F}_i (described later),

our objective is

$$\{\mathbf{W}_1, \dots, \mathbf{W}_N\} = \arg \min \mathcal{G}(\mathcal{F}_1, \dots, \mathcal{F}_N). \quad (1)$$

Typically, $\mathcal{G}(\mathcal{F}_1, \dots, \mathcal{F}_N) = \sum_{i=1}^N n_i \mathcal{F}_i$, $n_i = |\mathcal{D}_i| / \sum_{j=1}^N |\mathcal{D}_j|$, and $|\mathcal{D}_i|$ is the sample amount on client i .

3.2 Federated Conditional Policy (FedCP)

We focus on feature information separation for the feature vector

$$\mathbf{h}_i = f(\mathbf{x}_i; \mathbf{W}_i^{fe}), \forall (\mathbf{x}_i, y_i) \in \mathcal{D}_i. \quad (2)$$

Due to statistical heterogeneity, $\mathbf{h}_i \in \mathbb{R}^K$ contains global and personalized feature information. To separately exploit these two kinds of information, we propose **FedCP** that learns sample-specific separation in an end-to-end manner, as shown in Figure 2.

3.2.1 Separating feature information. Guided by the global information in the frozen global head and the personalized information in the personalized head, the CPN (the core of FedCP) can learn to generate the sample-specific policy and separate the global and personalized information in \mathbf{h}_i automatically.

Specifically, we devise CPN as the concatenation of an FC layer and a layer-normalization layer [3] followed by the ReLU activation function [28], as shown in Figure 2(a). On client i , we generate the sample-specific policy by

$$\{\mathbf{r}_i, \mathbf{s}_i\} := \text{CPN}(C_i; \Theta_i), \quad (3)$$

where $\mathbf{r}_i \in \mathbb{R}^K$, $\mathbf{s}_i \in \mathbb{R}^K$, $r_i^k + s_i^k = 1, \forall k \in [K]$, and $C_i \in \mathbb{R}^K$ is the sample-specific input for CPN. We describe the details of the input C_i and the output $\{\mathbf{r}_i, \mathbf{s}_i\}$ as follows.

C_i is generated to achieve the sample-specific characteristic and introduce personalized (client-specific) information. We can directly obtain the sample-specific vector \mathbf{h}_i , so we only introduce how to obtain the client-specific information here. Based on FedRep and FedRoD, the parameters in the personalized head, *i.e.*, \mathbf{W}_i^{hd} , naturally contain client-specific information. However, \mathbf{W}_i^{hd} is a matrix, not a vector. Thus, we generate \mathbf{v}_i by reducing the dimension of \mathbf{W}_i^{hd} . Recall that a head is an FC layer in FedCP, *i.e.*, $\mathbf{W}_i^{hd} \in \mathbb{R}^{C \times K}$, so the k th column of \mathbf{W}_i^{hd} corresponds to k th feature in \mathbf{h}_i . We obtain $\mathbf{v}_i := \sum_{c=1}^C \mathbf{w}_c^T$, where \mathbf{w}_c is the c th row in \mathbf{W}_i^{hd} and $\mathbf{v}_i \in \mathbb{R}^K$. In this way, we obtain a client-specific vector with the same shape and feature-wise semantics as \mathbf{h}_i . Then we combine sample-specific \mathbf{h}_i and the client-specific \mathbf{v}_i via $C_i := (\mathbf{v}_i / \|\mathbf{v}_i\|_2) \odot \mathbf{h}_i$, where $\|\mathbf{v}_i\|_2$ is the ℓ_2 -norm [36] of \mathbf{v}_i and \odot is the Hadamard product. We obtain \mathbf{v}_i before local learning in each iteration and regard it as a constant during training. During inference, we reuse the latest \mathbf{v}_i .

We separate information by multiplying the policy $\{\mathbf{r}_i, \mathbf{s}_i\}$ and \mathbf{h}_i to obtain the global feature information $\mathbf{r}_i \odot \mathbf{h}_i$ and personalized feature information $\mathbf{s}_i \odot \mathbf{h}_i$. There are connections among features [52], so we output $\{\mathbf{r}_i, \mathbf{s}_i\}$ with real numbers instead of Boolean values, *i.e.*, $r_i^k \in (0, 1)$ and $s_i^k \in (0, 1)$. Inspired by the Gumbel-Max trick for policy generating [11], we generate the policy with the help of the intermediates and a softmax [14] operation through the following two steps. Firstly, CPN generates the intermediates $\mathbf{a}_i \in \mathbb{R}^{K \times 2}$, where $\mathbf{a}_i^k = \{a_{i,1}^k, a_{i,2}^k\}, k \in [K]$, $a_{i,1}^k$ and $a_{i,2}^k$ are scalars without

Algorithm 1 The Learning Process in FedCP

Input: N clients with their local data, $\mathbf{W}^{fe,0}$: initial parameters of the global feature extractor, $\mathbf{W}^{hd,0}$: initial parameters of the global head, Θ^0 : initial parameters of the global CPN, η : local learning rate, λ : hyper-parameter for MMD loss, $\rho \in (0, 1]$: client joining ratio in one iteration, T : total training iterations.

Output: Reasonable personalized models $\{\mathbf{W}_1, \dots, \mathbf{W}_N\}$.

- 1: Server sends $\mathbf{W}^{fe,0}$ and $\mathbf{W}^{hd,0}$ to initialize $\mathbf{W}^{fe}, \mathbf{W}^{hd}, \mathbf{W}_i^{fe}$, and \mathbf{W}_i^{hd} on client $i, \forall i \in [N]$.
 - 2: Server sends Θ^0 to initialize the CPN on client $i, \forall i \in [N]$.
 - 3: **for** iteration $t = 0, \dots, T$ **do**
 - 4: Server randomly samples a subset \mathcal{I}^t of clients based on ρ .
 - 5: Server sends $\mathbf{W}^{fe,t}, \mathbf{W}^{hd,t}$, and Θ^t to the selected clients.
 - 6: **for** Client $i \in \mathcal{I}^t$ in parallel **do**
 - 7:

▷ **local initialization**

 Client i overwrites \mathbf{W}^{fe} and \mathbf{W}_i^{fe} with the parameters $\mathbf{W}^{fe,t}$ and freezes \mathbf{W}^{fe} .
 - 8: Client i overwrites \mathbf{W}^{hd} with the parameters $\mathbf{W}^{hd,t}$ and freezes \mathbf{W}^{hd} .
 - 9: Client i overwrites Θ_i with the parameters Θ^t .
 - 10: Client i generates the client-specific vector \mathbf{v}_i .
 - 11:

▷ **local learning**

 Client i updates $\mathbf{W}_i^{fe}, \mathbf{W}_i^{hd}$ and Θ_i simultaneously:
 - 12: $\mathbf{W}_i^{fe} \leftarrow \mathbf{W}_i^{fe} - \eta \nabla_{\mathbf{W}_i^{fe}} \mathcal{F}_i$;
 - 13: $\mathbf{W}_i^{hd} \leftarrow \mathbf{W}_i^{hd} - \eta \nabla_{\mathbf{W}_i^{hd}} \mathcal{F}_i$;
 - 14: $\Theta_i \leftarrow \Theta_i - \eta \nabla_{\Theta_i} \mathcal{F}_i$.
 - 15: Client i obtains $\widehat{\mathbf{W}}_i^{hd}$ through Eq. (6).
 - 16: Client i uploads $\{\mathbf{W}_i^{fe}, \widehat{\mathbf{W}}_i^{hd}, \Theta_i\}$ to the server.
 - 17: **end for**
 - 18:

▷ **Server aggregation**

 Server calculates $n^t = \sum_{i \in \mathcal{I}^t} n_i$ and obtains
 - 19: $\mathbf{W}^{fe,t+1} = \frac{1}{n^t} \sum_{i \in \mathcal{I}^t} n_i \mathbf{W}_i^{fe}$;
 - 20: $\mathbf{W}^{hd,t+1} = \frac{1}{n^t} \sum_{i \in \mathcal{I}^t} n_i \widehat{\mathbf{W}}_i^{hd}$;
 - 21: $\Theta^{t+1} = \frac{1}{n^t} \sum_{i \in \mathcal{I}^t} n_i \Theta_i$.
 - 22: **end for**
 - 23: **return** $\{\mathbf{W}_1, \dots, \mathbf{W}_N\}$
-

constraint. Secondly, we obtain r_i^k and s_i^k by

$$r_i^k = \frac{\exp(a_{i,1}^k)}{\sum_{j \in \{1,2\}} \exp(a_{i,j}^k)}, \quad s_i^k = \frac{\exp(a_{i,2}^k)}{\sum_{j \in \{1,2\}} \exp(a_{i,j}^k)}. \quad (4)$$

Note that, $r_i^k \in (0, 1), s_i^k \in (0, 1), r_i^k + s_i^k = 1, \forall k \in [K]$ still holds.

3.2.2 Processing feature information. Then, we feed $\mathbf{r}_i \odot \mathbf{h}_i$ and $\mathbf{s}_i \odot \mathbf{h}_i$ to the global head and the personalized head, respectively. The outputs of global head and the personalized head are $\text{out}_i^r = g(\mathbf{r}_i \odot \mathbf{h}_i; \mathbf{W}^{hd})$ and $\text{out}_i^s = g(\mathbf{s}_i \odot \mathbf{h}_i; \mathbf{W}_i^{hd})$, respectively. We define the final output $\text{out}_i := \text{out}_i^r + \text{out}_i^s$. Then the local loss is

$$\mathcal{E}_i = \mathbb{E}_{(\mathbf{x}_i, y_i) \sim \mathcal{D}_i} \mathcal{L}(\text{out}_i, y_i), \quad (5)$$

where \mathcal{L} is the cross-entropy loss function [33].

From the view of each sample, the extracted features are processed by both the global head and the personalized head. For simplicity, we aggregate these two heads through averaging to form the upload head \widehat{W}_i^{hd} :

$$\widehat{W}_i^{hd} = \frac{W^{hd} + W_i^{hd}}{2}. \quad (6)$$

In each iteration, we upload $\{W_i^{fe}, \widehat{W}_i^{hd}, \Theta_i\}$ to the server.

3.2.3 Aligning features. To fit the features outputted by the personalized feature extractor with the frozen global head, we align the features outputted by the personalized feature extractor and the global feature extractor through the MMD loss \mathcal{E}_i^d ,

$$\mathcal{E}_i^d = \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_i} \kappa[\mathbf{h}_i, f(\mathbf{x}_i; W^{fe})], \quad (7)$$

where κ is the radial basis function (RBF) kernel. Finally, we have the local loss $\mathcal{F}_i = \mathcal{E}_i + \lambda \mathcal{E}_i^d$, where λ is a hyper-parameter. Specifically,

$$\mathcal{F}_i = \mathbb{E}_{(x_i, y_i) \sim \mathcal{D}_i} \{ \mathcal{L}[g(\mathbf{r}_i \circ \mathbf{h}_i; W^{hd}) + g(\mathbf{s}_i \circ \mathbf{h}_i; W_i^{hd}), y_i] + \lambda \kappa[\mathbf{h}_i, f(\mathbf{x}_i; W^{fe})] \}, \quad (8)$$

where \mathbf{h}_i is the feature vector extracted by Eq. (2), and \mathbf{r}_i and \mathbf{s}_i are obtained through Eq. (3).

We show the entire *learning* process in Algorithm 1. For *inference*, we use the personalized model as illustrated by the non-gray border modules in Figure 2(a).

3.3 Privacy Analysis

According to Figure 2(b) and Algorithm 1, our proposed FedCP shares the parameters of one feature extractor, one head, and one CPN. As for the head part, we upload \widehat{W}_i^{hd} on each client after aggregating W^{hd} and W_i^{hd} by Eq. (6). This process can be viewed as adding noise (global parameters W^{hd}) to W_i^{hd} , thus protecting privacy during the uploading and downloading. Besides, the sample-specific characteristic further improves the privacy-preserving ability of FedCP. On the one hand, since C_i is dynamically generated without sharing with the server, it is hard to recover the sample-specific policy with the CPN or through model inversion attacks [1]. On the other hand, without the sample-specific policy, the connection between the feature extractor and the head is broken, increasing the difficulty of attacks based on shared model parameters. We evaluate the privacy-preserving ability of FedCP in Appendix B.

4 EXPERIMENTAL SETUP

We evaluate FedCP on various image/text classification tasks. For the image classification tasks, we use four famous datasets, including MNIST [22], Cifar10 [20], Cifar100 [20] and Tiny-ImageNet [5] (100K images with 200 classes) using a famous 4-layer CNN [9, 31, 32]. To evaluate FedCP on a larger backbone model than the 4-layer CNN, we also use ResNet-18 [13] on Tiny-ImageNet. We set the local learning rate $\eta = 0.005$ for the 4-layer CNN and $\eta = 0.1$ for ResNet-18. For the text classification tasks, we use the AG News [58] dataset with the fastText [18] and set $\eta = 0.1$ for fastText with other settings being the same as image classification tasks.

We simulate the heterogeneous settings in two widely-used scenarios, *i.e.*, the pathological setting [32, 41] and practical setting [23, 29]. For the pathological setting, we sample 2/2/10 classes

on MNIST/Cifar10/Cifar100 from a total of 10/10/100 classes for each client with disjoint data. Specifically, similar to FedAvg [32], we separate clients into groups that own unbalanced data with the same labels. Following MOON [23], we create the practical setting through the Dirichlet distribution, denoted as $Dir(\beta)$. Specifically, we sample $q_{c,i} \sim Dir(\beta)$ and allocate a $q_{c,i}$ proportion of the samples of class c to client i . We set $\beta = 0.1$ for the default practical setting [29, 46]. Then, we split the data on each client into a training dataset (75%) and a test dataset (25%).

Following FedAvg, we set the local batch size to 10 and the number of local learning epochs to 1. We run all tasks up to 2000 iterations until all methods converge empirically. Based on pFedMe, FedFomo, and FedRoD, we set the total number of clients to 20 and the client joining ratio $\rho = 1$ by default. Following pFedMe, we report the test accuracy of the best global model for traditional FL methods and the average test accuracy of the best personalized models for pFL methods. We run all the experiments five times and report the mean and standard deviation. Besides, we run all experiments on a machine with two Intel Xeon Gold 6140 CPUs (36 cores), 128G memory, eight NVIDIA 2080 Ti GPUs, and CentOS 7.8. For more results and details, please refer to the Appendix.




Figure 3: The first row shows six samples from Tiny-ImageNet. The second and third rows respectively show the Grad-CAM visualizations of the learned personalized model with only the global head or the personalized head activated. Highlighted areas are the parts the model pays attention to.

5 ABLATION STUDY

5.1 Feature Information Visualization

To visualize the separated global and personalized feature information when using ResNet-18, we adopt the Grad-CAM [40] on the learned personalized model when only the global head or the personalized head is activated. Six cases from Tiny-ImageNet are shown in Figure 3.

According to Figure 3, with only the global head activated, the personalized model focuses on relatively global information, such as trees (Case 0 and Case 4), grasses (Case 1), or sky (Case 2 and Case 5) in the background. When we only activate the personalized head, the personalized model focuses on the relatively personalized information, such as foreground (Case 2 and Case 5) or objects (Case 0, Case 1, and Case 4). As for Case 3, the rarely-used pink color is more personalized than the widely-used blue color.




Figure 4: Illustration of variants for module ablation study.

Table 1: The accuracy (%) on Tiny-ImageNet using ResNet-18 for ablation study.

FedCP	w.o. cs	w.o. ss	w.o. cs & ss	w.o. GFM	w.o. CPN	w.o. CPN & GFM	w.o. CPN & GH	w.o. CPN & GFM & GH
44.18±0.21	43.76±0.39	42.73±0.26	42.25±0.30	42.87±0.36	41.17±0.18	40.06±0.47	35.44±0.78	39.04±0.79

5.2 Effectiveness of CPN input

To show the effectiveness of each part of the CPN input, we remove them one by one and obtain the variants: without client-specific vector (*w.o. cs*), without sample-specific vector (*w.o. ss*), without client-specific and sample-specific vector (*w.o. cs & ss*). For *w.o. cs & ss*, we regard the randomly initialized frozen vector as the CPN input, which has the same shape as the sample-specific vector.

In Table 1, removing either the client-specific vector or the sample-specific vector causes an accuracy decrease. However, *w.o. cs* performs better than *w.o. ss*, so the sample-specific vector is more significant than the client-specific one. According to Table 1 and Table 2, removing these two kinds of information and using the random vector, *w.o. cs & ss* still achieves higher accuracy than all the baselines because CPN module can still learn to separate feature information through the end-to-end training.

5.3 Effectiveness of FedCP modules

To show the effectiveness of each module in FedCP, we remove them one by one and obtain the variants: without the frozen global feature extractor and the MMD loss (without GFM for short, *i.e.*, *w.o. GFM*), without CPN (*w.o. CPN*), without CPN and GFM (*w.o. CPN & GFM*), without CPN and the frozen global head (*w.o. CPN & GH*), without CPN, GFM, and the frozen global head (*w.o. CPN & GFM & GH*, similar to FedPer), as shown in Figure 4. It is invalid to keep CPN while removing the frozen global head since they are a **union** for our feature separating goal.

In Table 1, without the GFM to align the features, the accuracy of *w.o. GFM* decreases by 1.31% compared to FedCP, but it still outperforms other baselines (see Table 2). Without CPN, the accuracy of *w.o. CPN* decreases by 3.01%, so CPN is more critical than the GFM when the frozen global head exists. Removing both the CPN and the GFM (*w.o. CPN & GFM*) degenerates further than removing one of them, which means that these two modules can facilitate each other. The CPN and the frozen global head are the key modules in FedCP. Without them, the performance of *w.o. CPN & GH* degenerates significantly, with a 8.74% drop compared to FedCP. Furthermore, *w.o. CPN & GFM & GH* (removing all the modules) performs better than *w.o. CPN & GH*. It means simply adding the GFM to *w.o. CPN & GH* causes performance degeneration.

6 EVALUATION AND ANALYSIS

6.1 Main Experiments

Due to the limited space, we use the “TINY” and “TINY*” to represent using the 4-layer CNN on Tiny-ImageNet and using ResNet-18 on Tiny-ImageNet, respectively. Table 2 shows that FedCP outperforms all the baselines when using either the 4-layer CNN or the ResNet-18, especially on relatively challenging tasks. In the default practical setting on Cifar100, FedCP exceeds the best baseline (Ditto) by **6.69%**. Our CPN only introduces an additional 0.527M (million) parameters on each client, which is 9.25% and 4.67% of the parameters in the 4-layer CNN (5.695M) and the ResNet-18 (11.279M), respectively. In the following, we analyze **why** FedCP outperforms all the baselines.

In Table 2, FedAvg and FedProx perform poorly, as the global model cannot fit the local data well on all the clients. They directly feed features to the global head, regardless of the personalized information in the features. In contrast, FedCP separates and feeds the global information and the personalized information in the features to the global head and the personalized head, respectively.

Per-FedAvg performs poorly among pFL methods, as the aggregated learning trend can hardly meet the trend of each personalized model. In contrast, FedCP considers personalization in a sample-specific manner conditioned by the client-specific vector, which meets the demand of each client, thus performing better.

pFedMe and FedAMP utilize regularization terms to extract information from the local model and the client-specific server model, respectively. However, excessively concentrating on personalization is not beneficial to the collaborative goal of FL. Since Ditto extracts global information from the global model, it performs better than pFedMe and FedAMP. Like Ditto, FedCP also takes advantage of global information for each client.

FedPer and FedRep only share the feature extractor without sharing heads. They ignore some global information in the head part, so they perform worse than FedCP. FedRoD bridges the goal of traditional FL and pFL by learning two heads with two objectives. However, these two goals are competing [4], so FedRoD performs worse than FedRep, which also learns a personalized head but only focuses on the goal of pFL. Like FedRep, FedCP only focuses on the pFL goal, thus performing the best.

Table 2: The accuracy (%) of the image/text classification tasks in the main experiments.

Settings	Pathological setting			Default practical setting ($\beta = 0.1$)					
	MNIST	Cifar10	Cifar100	MNIST	Cifar10	Cifar100	TINY	TINY*	AG News
FedAvg [32]	97.93±0.05	55.09±0.83	25.98±0.13	98.81±0.01	59.16±0.47	31.89±0.47	19.46±0.20	19.45±0.13	79.57±0.17
FedProx [26]	98.01±0.09	55.06±0.75	25.94±0.16	98.82±0.01	59.21±0.40	31.99±0.41	19.37±0.22	19.27±0.23	79.35±0.23
Per-FedAvg [8]	99.63±0.02	89.63±0.23	56.80±0.26	98.90±0.05	87.74±0.19	44.28±0.33	25.07±0.07	21.81±0.54	93.27±0.25
pFedMe [42]	99.75±0.02	90.11±0.10	58.20±0.14	99.52±0.02	88.09±0.32	47.34±0.46	26.93±0.19	33.44±0.33	91.41±0.22
FedAMP [15]	99.76±0.02	90.79±0.16	64.34±0.37	99.47±0.02	88.70±0.18	47.69±0.49	27.99±0.11	29.11±0.15	94.18±0.09
Ditto [24]	99.81±0.00	92.39±0.06	67.23±0.07	99.64±0.00	90.59±0.01	52.87±0.64	32.15±0.04	35.92±0.43	95.45±0.17
FedPer [2]	99.70±0.02	91.15±0.21	63.53±0.21	99.47±0.04	89.22±0.33	49.63±0.54	33.84±0.34	38.45±0.85	95.54±0.32
FedRep [6]	99.77±0.03	91.93±0.14	67.56±0.31	99.48±0.02	90.40±0.24	52.39±0.35	37.27±0.20	39.95±0.61	96.28±0.14
FedRoD [4]	99.90±0.00	91.98±0.03	62.30±0.02	99.66±0.00	89.93±0.01	50.94±0.11	36.43±0.05	37.99±0.26	95.99±0.08
FedFomo [56]	99.83±0.00	91.85±0.02	62.49±0.22	99.33±0.04	88.06±0.02	45.39±0.45	26.33±0.22	26.84±0.11	95.84±0.15
FedPHP [27]	99.73±0.00	90.01±0.00	63.09±0.04	99.58±0.00	88.92±0.02	50.52±0.16	35.69±3.26	29.90±0.51	94.38±0.12
FedCP	99.91±0.01	92.67±0.09	71.80±0.16	99.71±0.00	91.30±0.17	59.56±0.08	43.49±0.04	44.18±0.21	96.78±0.09

Similar to FedAMP, FedFomo aggregates client models with client-specific weights, thus losing some global information. FedPHP transfers the global information only in the global feature extractor through the MMD loss. Although it achieves excellent performance, FedPHP loses the global information in the global head during local training, so it performs worse than FedCP.

6.2 Computing and Communication Overhead

Here, we focus on the training phase. We report the total time and the number of iterations required for each method to converge and calculate the average time consumption in each iteration, as shown in Table 3. Ditto and pFedMe cost more time in each iteration than most methods since the additional personalized model training takes much extra time. Compared to most baselines, *e.g.*, Per-FedAvg, pFedMe, Ditto, FedRep, and FedPHP, FedCP costs less training time in each iteration. In FedCP, the parameters in the CPN module only require an additional 4.67% communication overhead per iteration when using ResNet-18 compared to FedAvg.

Table 3: The computing time and communication iterations on Tiny-ImageNet using ResNet-18.

	Total time	Iterations	Avg. time
FedAvg	365 min	230	1.59 min
FedProx	325 min	163	1.99 min
Per-FedAvg	121 min	34	3.56 min
pFedMe	1157 min	113	10.24 min
FedAMP	92 min	60	1.53 min
Ditto	318 min	27	11.78 min
FedPer	83 min	43	1.92 min
FedRep	471 min	115	4.09 min
FedRoD	87 min	50	1.74 min
FedFomo	193 min	71	2.72 min
FedPHP	264 min	65	4.06 min
FedCP	204 min	74	2.75 min

6.3 Different Heterogeneity Degrees

In addition to Table 2, we conduct experiments on the settings with different degrees of heterogeneity on Tiny-ImageNet and AG News by varying β . The smaller the β is, the more heterogeneous the setting is. We show the accuracy in Table 4, where FedCP still outperforms the baselines. Most pFL methods achieve higher accuracy than traditional FL methods in the more heterogeneous setting. In the setting with a larger β , most of them cannot achieve higher accuracy than FedAvg on Tiny-ImageNet. In contrast, the methods that utilize global information during local learning (FedPHP, FedRoD, and FedCP) maintain excellent performance. FedRoD performs worse than FedRep, as the latter focuses only on the goal of pFL. pFedMe and FedAMP perform poorly among pFL methods. Their accuracy is lower than traditional FL methods when $\beta = 1$.

6.4 Scalability with Different Client Amounts

Following MOON [23], we conduct another six experiments (*i.e.*, $N = 10$, $N = 30$, $N = 50$, $N = 100$, $N = 200$, and $N = 500$) to study the scalability of FedCP and keep other settings unchanged. Per-FedAvg requires more data than other methods, as meta-learning requires at least two batches of data, which is invalid on some clients in our unbalanced settings when $N \geq 200$. Since the total data amount is constant on Cifar100, the local data amount (on average) decreases as the client amount increases. With both N and local data amount changing, it is unreasonable to compare the results among different N in Table 4. Some pFL methods, including Per-FedAvg and pFedMe, achieve relatively poor performance in the setting with $N = 10$, where few clients (*e.g.*, hospitals) participate in FL, and each of them possesses a large data repository. When $N = 500$ (*e.g.*, mobile smartphones), each client only has 90 samples for training on average, which is not enough for the weight calculation in FedFomo, so it performs worse than FedAvg. FedAMP diverges as it is hard to find similar clients when they have little data. According to Table 4, FedCP still outperforms all the baselines.

To simulate a real-world scenario where more clients means more total data amount in FL, we consider the setting Cifar100 ($\beta = 0.1$, $\rho = 1$, and $N = 50$) used above as the base setting and

Table 4: The accuracy (%) of the image/text classification tasks for heterogeneity and scalability.

Datasets	Heterogeneity			Scalability					
	TINY		AG News	Cifar100					
	$\beta = 0.01$	$\beta = 0.5$	$\beta = 1$	$N = 10$	$N = 30$	$N = 50$	$N = 100$	$N = 200$	$N = 500$
FedAvg	15.70±0.46	21.14±0.47	87.12±0.19	31.47±0.01	31.15±0.05	31.90±0.27	31.95±0.37	31.20±0.58	29.51±0.73
FedProx	15.66±0.36	21.22±0.47	87.21±0.13	31.24±0.08	31.21±0.08	31.94±0.30	31.97±0.24	31.22±0.62	29.84±0.81
Per-FedAvg	39.39±0.30	16.36±0.13	87.08±0.26	37.24±0.12	41.57±0.21	44.31±0.20	36.07±0.24	—	—
pFedMe	41.45±0.14	17.48±0.61	87.08±0.18	44.06±0.29	47.04±0.28	48.36±0.64	46.45±0.18	39.55±0.61	31.30±0.89
FedAMP	48.42±0.06	12.48±0.21	83.35±0.05	49.23±0.18	45.33±0.04	44.39±0.35	40.43±0.17	35.40±0.70	<i>diverged</i>
Ditto	50.62±0.02	18.98±0.05	91.89±0.17	52.32±0.19	52.53±0.42	54.22±0.04	52.89±0.22	35.18±0.53	30.24±0.72
FedPer	51.83±0.22	17.31±0.19	91.85±0.24	50.31±0.19	44.98±0.20	44.22±0.18	40.37±0.41	34.99±0.48	30.56±0.59
FedRep	55.43±0.15	16.74±0.09	92.25±0.20	52.89±0.10	50.24±0.01	47.41±0.18	44.61±0.20	36.79±0.60	31.92±0.71
FedRoD	49.17±0.06	23.23±0.11	92.16±0.12	49.83±0.07	50.11±0.03	49.38±0.01	46.65±0.22	43.53±0.86	34.61±0.98
FedFomo	46.36±0.54	11.59±0.11	91.20±0.18	46.71±0.23	43.20±0.05	42.56±0.33	38.91±0.08	34.79±0.71	29.24±1.28
FedPHP	48.63±0.02	21.09±0.07	90.52±0.19	49.32±0.19	49.28±0.06	52.44±0.16	49.70±0.31	34.48±0.33	30.26±0.84
FedCP	56.31±0.39	27.66±0.16	92.89±0.10	58.36±0.02	56.93±0.19	55.43±0.21	53.81±0.32	44.86±0.87	35.87±0.52

Table 5: The accuracy (%) on Cifar100 for scalability.

	$N = 10 50$	$N = 30 50$	$N = 50$
FedAvg	25.28±0.32	29.04±0.21	31.90±0.27
FedProx	25.65±0.34	29.04±0.36	31.94±0.30
Per-FedAvg	40.20±0.21	42.96±0.42	44.31±0.20
pFedMe	40.27±0.54	42.19±0.38	48.36±0.64
FedAMP	43.57±0.30	43.18±0.31	44.39±0.35
Ditto	48.23±0.35	50.98±0.29	54.22±0.04
FedPer	43.64±0.42	43.54±0.43	44.22±0.18
FedRep	46.85±0.12	47.63±0.26	47.41±0.18
FedRoD	46.32±0.02	49.15±0.12	49.38±0.01
FedFomo	41.53±0.45	40.69±0.41	42.56±0.33
FedPHP	45.71±0.21	48.65±0.24	52.44±0.16
FedCP	50.93±0.34	54.31±0.25	55.43±0.21

randomly sample 10 and 30 clients from existing 50 clients to form the Cifar100 ($\beta = 0.1, \rho = 1$, and $N = 10|50$) and Cifar100 ($\beta = 0.1, \rho = 1$, and $N = 30|50$) settings, respectively. When we increase the client amount, the accuracy increases as more data are utilized to train the globally shared modules, which facilitates information transfer among clients. The superior performance of FedCP in Table 5 shows its scalability in this real-world scenario.

6.5 Large Local Epochs

Large local epochs can reduce total communication iterations but increase computing overhead per iteration for most of the methods in FL [32]. With larger local epochs, FedCP can still maintain its superiority as shown in Table 6. Most of the methods perform worse with larger local epochs since more local training aggravates the discrepancy among client models, which is adverse to server aggregation. For example, the accuracy of FedRoD drops by 2.16% when the number of local epochs increases from 5 to 40.

Table 6: The accuracy (%) on Cifar10 in the default practical setting with large local epochs.

Local epochs	5	10	20	40
FedAvg	57.51±0.35	57.55±0.32	57.28±0.23	56.27±0.29
FedProx	57.48±0.28	57.69±0.31	57.53±0.33	56.18±0.24
Per-FedAvg	86.13±0.12	86.09±0.19	85.57±0.15	85.45±0.16
pFedMe	88.72±0.02	88.58±0.17	88.37±0.14	88.16±0.20
FedAMP	88.72±0.21	88.77±0.27	88.76±0.30	88.70±0.26
Ditto	90.79±0.21	90.59±0.06	90.34±0.23	90.02±0.38
FedPer	89.62±0.12	89.73±0.31	89.79±0.35	89.49±0.55
FedRep	90.20±0.41	90.08±0.26	89.46±0.13	89.22±0.25
FedRoD	89.71±0.32	89.11±0.33	88.13±0.21	87.55±0.28
FedFomo	88.39±0.15	88.43±0.16	88.41±0.13	88.13±0.32
FedPHP	90.29±0.37	90.03±0.23	89.92±0.27	89.87±0.26
FedCP	91.13±0.34	91.24±0.31	91.02±0.28	90.86±0.37

6.6 Clients Accidentally Dropping Out

Table 7: The accuracy (%) on Cifar100 ($N = 50, \beta = 0.1$) when clients accidentally drop out.

	$\rho = 1$	$\rho \in [0.5, 1]$	$\rho \in [0.1, 1]$
Per-FedAvg	44.31±0.20	43.66±1.38	43.63±1.07
pFedMe	48.36±0.64	43.28±0.85	41.71±1.02
FedAMP	44.39±0.35	42.91±0.08	42.92±0.14
Ditto	50.59±0.22	49.78±0.36	48.33±3.27
FedPer	44.22±0.18	44.12±0.21	44.07±0.27
FedRep	47.41±0.18	46.93±0.21	46.61±0.22
FedRoD	49.38±0.01	49.07±0.43	47.80±1.35
FedFomo	42.56±0.33	40.96±0.02	40.93±0.07
FedPHP	50.23±0.12	45.19±0.07	44.43±0.12
FedCP	54.81±0.20	54.68±0.35	54.20±0.21

Due to the changing network connection quality, some clients may accidentally (randomly) drop out at one iteration and become active again at another iteration, which frequently happens in the mobile settings. We compare the performance of pFL methods when some clients accidentally drop out, as shown in Table 7. Instead of using the constant ρ , we randomly choose a value within a given range for ρ in each iteration. The larger the range of ρ is, the more unstable the setting is. It simulates a more practical setting with a random drop-out rate than the settings used by the SOTA methods, which set a constant drop-out rate in all iterations.

Most pFL methods suffer from an accuracy decrease in unstable settings. pFedMe and FedPHP have up to 6.65% and 9.80% accuracy decrease, respectively, compared to $\rho = 1$ in Table 7. Some methods, such as FedRep, and FedRoD, perform worse with a larger range of ρ . The standard deviation of Per-FedAvg, pFedMe, Ditto, and FedRoD is greater than 1% when $\rho \in [0.1, 1]$, which means their performance is unstable with the random ρ . Since CPN separates feature information automatically, FedCP can adapt to the changing environments thus still maintaining superiority and stable performance in these unstable settings.

7 EFFECT OF THE HYPER-PARAMETER λ

To guide the learned features to fit the frozen global head, we use the hyper-parameter λ to control the importance of MMD loss that aligns the outputs of the personalized feature extractor and the outputs of the global feature extractor. The larger the λ is, the closer these two outputs are.

Table 8: The accuracy (%) on Tiny-ImageNet using the 4-layer CNN in three practical settings.

	$\lambda = 1$	$\lambda = 2$	$\lambda = 5$	$\lambda = 10$	$\lambda = 50$
$\beta = 0.01$	56.56±0.35	56.71±0.32	56.31±0.39	54.48±0.10	9.73±0.02
$\beta = 0.1$	41.67±0.17	42.75±0.03	43.49±0.04	42.83±0.07	8.14±0.06
$\beta = 0.5$	24.95±0.15	26.55±0.23	27.66±0.16	26.95±0.27	4.54±0.04

From Table 8, the accuracy first increases and then decreases as λ increases, which is similar among three settings with different degrees of heterogeneity. By assigning a proper value to λ , the personalized feature extractor can learn the information from the local data while guiding the output features to fit the frozen global head. When the value of λ is overlarge (e.g., $\lambda = 50$), the personalized feature extractor can hardly learn from the local data. Instead, it tends to output similarly to the frozen global feature extractor. To pay more attention to the local data in a more heterogeneous setting (e.g., $\beta = 0.01$), FedCP requires a relatively smaller λ , as the global information plays a less critical role in this situation.

8 POLICY STUDY

We show the policy change for the training samples and the generated policies for all the test samples during inference in Figure 5. For clarity, we collect all the sample-specific s_i on each client and average them to obtain \bar{s}_i . Then we further average the elements in \bar{s}_i to generate one scalar, which is called personalization identification ratio (PIR): $\text{PIR}_i := \frac{1}{K} \sum_k \bar{s}_i^k$, $i \in [N]$, where \bar{s}_i^k is the k th element in the policy \bar{s}_i .




Figure 5: Visualizations for PIR and s_i distribution on Tiny-ImageNet in the default practical setting. Blue color and orange color represent the figures for the 4-layer CNN and ResNet-18, respectively. We draw PIR change curves for training samples. Best viewed in color.

When using diverse backbones with different feature extraction abilities, the policies vary in both PIR change and s_i distribution. As shown in Figure 5(a), on client #0, PIR increases from the initial value of 0.50 to around 0.58 in the first 20 iterations and remains almost unchanged using the 4-layer CNN. However, when using ResNet-18, PIR decreases first and then increases rapidly to around 0.61, which means that the features extracted by the feature extractor in ResNet-18 contain more global feature information in early iterations, and our CPN can automatically capture this dynamic characteristic during all FL iterations. In Figure 5(b), the value range of s_i varies among clients, as they contain diverse samples. For example, the s_i range on client #10 is the largest among clients. Although the policies are different for the samples, the mean values of s_i are similar among clients when using one specific backbone, as shown in Figure 5(b). The values of s_i are all larger than 0.5 during inference, which means the learned features contain more personalized feature information than global feature information on clients in these scenarios.

9 CONCLUSION

We propose a Federated Conditional Policy (FedCP) method that generates a policy for each sample to separate its features into the global feature information and the personalized feature information, then processes them by the global head and the personalized head, respectively. FedCP outperforms eleven SOTA methods by up to 6.69% under various settings with excellent privacy-preserving ability. Besides, FedCP also maintains excellent performance when some clients accidentally drop out.

ACKNOWLEDGMENTS

This work was supported in part by the Shanghai Key Laboratory of Scalable Computing and Systems, National Key R&D Program of China (2022YFB4402102), Internet of Things special subject program, China Institute of IoT (Wuxi), Wuxi IoT Innovation Promotion Center (2022SP-T13-C), Industry-university-research Cooperation Funding Project from the Eighth Research Institute in China Aerospace Science and Technology Corporation (Shanghai) (USCAST2022-17), and Intel Corporation (UFunding 12679). The work of H. Wang was supported in part by the NSF grant CRII-OAC-2153502. Ruhui Ma is the corresponding author.

REFERENCES

- [1] Mohammad Al-Rubaie and J Morris Chang. 2016. Reconstruction attacks against mobile-based continuous authentication systems in the cloud. *IEEE Transactions on Information Forensics and Security* 11, 12 (2016), 2648–2663.
- [2] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. 2019. Federated Learning with Personalization Layers. *arXiv preprint arXiv:1912.00818* (2019).
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [4] Hong-You Chen and Wei-Lun Chao. 2021. On Bridging Generic and Personalized Federated Learning for Image Classification. In *International Conference on Learning Representations (ICLR)*.
- [5] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. 2017. A Downsampled Variant of Imagenet as an Alternative to the Cifar Datasets. *arXiv preprint arXiv:1707.08819* (2017).
- [6] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. 2021. Exploiting Shared Representations for Personalized Federated Learning. In *International Conference on Machine Learning (ICML)*.
- [7] Lydia de la Torre. 2018. A guide to the california consumer privacy act of 2018. Available at SSRN 3275571 (2018).
- [8] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. 2020. Personalized Federated Learning with Theoretical Guarantees: A Model-Agnostic Meta-Learning Approach. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [9] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting gradients-how easy is it to break privacy in federated learning?. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [10] Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. 2006. A Kernel Method for the Two-Sample-Problem. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [11] Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, and Rogerio Feris. 2019. Spottune: Transfer Learning through Adaptive Fine-Tuning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [12] Seok-Ju Hahn, Minwoo Jeong, and Junhye Lee. 2022. Connecting Low-Loss Subspace for Personalized Federated Learning. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [14] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. Distilling the Knowledge in a Neural Network. *arXiv preprint arXiv:1503.02531* 2, 7 (2015).
- [15] Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. 2021. Personalized Cross-Silo Federated Learning on Non-IID Data. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [16] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *International Conference on Machine Learning (ICML)*.
- [17] Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical Reparameterization with Gumbel-Softmax. *arXiv preprint arXiv:1611.01144* (2016).
- [18] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of Tricks for Efficient Text Classification. In *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*.
- [19] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2019. Advances and Open Problems in Federated Learning. *arXiv preprint arXiv:1912.04977* (2019).
- [20] Alex Krizhevsky and Hinton Geoffrey. 2009. Learning Multiple Layers of Features From Tiny Images. *Technical Report* (2009).
- [21] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep Learning. *Nature* 521, 7553 (2015), 436–444.
- [22] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [23] Qinbin Li, Bingsheng He, and Dawn Song. 2021. Model-Contrastive Federated Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [24] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2021. Ditto: Fair and Robust Federated Learning Through Personalization. In *International Conference on Machine Learning (ICML)*.
- [25] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. 2020. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* 37, 3 (2020), 50–60.
- [26] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *Conference on Machine Learning and Systems (MLSys)*.
- [27] Xin-Chun Li, De-Chuan Zhan, Yunfeng Shao, Bingshuai Li, and Shaoming Song. 2021. FedPHP: Federated Personalization with Inherited Private Models. In *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML)*.
- [28] Yuanzhi Li and Yang Yuan. 2017. Convergence Analysis of Two-Layer Neural Networks with Relu Activation. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [29] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. 2020. Ensemble Distillation for Robust Model Fusion in Federated Learning. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [30] Lanlan Liu and Jia Deng. 2018. Dynamic Deep Neural Networks: Optimizing Accuracy-Efficiency Trade-Offs by Selective Execution. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [31] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. 2021. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [32] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*.
- [33] Kevin P Murphy. 2012. *Machine learning: a probabilistic perspective*. MIT press.
- [34] Dinh C Nguyen, Ming Ding, Pubudu N Pathirana, Aruna Seneviratne, Jun Li, and H Vincent Poor. 2021. Federated learning for internet of things: A comprehensive survey. *IEEE Communications Surveys & Tutorials* 23, 3 (2021), 1622–1658.
- [35] Boris Oreshkin, Pau Rodriguez López, and Alexandre Lacoste. 2018. Tadam: Task Dependent Adaptive Metric for Improved Few-Shot Learning. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [36] Florent Perronnin, Jorge Sánchez, and Thomas Mensink. 2010. Improving the Fisher Kernel for Large-scale Image Classification. In *European Conference on Computer Vision (ECCV)*.
- [37] Can Qin, Haoxuan You, Lichen Wang, C-C Jay Kuo, and Yun Fu. 2019. Pointdan: A Multi-Scale 3D Domain Adaption Network for Point Cloud Representation. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [38] Protection Regulation. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council. *Regulation (eu)* 679 (2016), 2016.
- [39] Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. 2020. Balanced Meta-softmax for Long-tailed Visual Recognition. *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [40] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual Explanations from Deep Networks via Gradient-based Localization. In *IEEE International Conference on Computer Vision (ICCV)*.
- [41] Aviv Shamsian, Aviv Navon, Ethan Fetaya, and Gal Chechik. 2021. Personalized Federated Learning using Hypernetworks. In *International Conference on Machine Learning (ICML)*.
- [42] Canh T Dinh, Nguyen Tran, and Tuan Dung Nguyen. 2020. Personalized Federated Learning with Moreau Envelopes. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [43] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. 2022. Towards Personalized Federated Learning. *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [44] Andreas Veit and Serge Belongie. 2018. Convolutional Networks with Adaptive Inference Graphs. In *European Conference on Computer Vision (ECCV)*.
- [45] Chi Wang, Yang Hua, Zheng Lu, Jian Gao, and Neil Robertson. 2021. Temporal Meta-Adaptor for Video Object Detection. In *British Machine Vision Conference (BMVC)*.
- [46] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. 2020. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [47] Yuxin Wu and Kaiming He. 2018. Group Normalization. In *European Conference on Computer Vision (ECCV)*.
- [48] Yuezhou Wu, Yan Kang, Jiahuan Luo, Yuanqin He, Lixin Fan, Rong Pan, and Qiang Yang. 2022. FedCG: Leverage Conditional GAN for Protecting Privacy and Maintaining Competitive Performance in Federated Learning. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [49] Liu Yang, Ben Tan, Vincent W Zheng, Kai Chen, and Qiang Yang. 2020. Federated Recommendation Systems. In *Federated Learning*. 225–239.
- [50] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated Machine Learning: Concept and Applications. *ACM Transactions on Intelligent Systems and Technology* 10, 2 (2019), 1–19.
- [51] Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, and Yanfeng Wang. 2023. FedDisco: Federated Learning with Discrepancy-Aware Collaboration. *arXiv preprint arXiv:2305.19229* (2023).
- [52] Lei Yu and Huan Liu. 2003. Feature Selection for High-Dimensional Data: A Fast Correlation-Based Filter Solution. In *International Conference on Machine Learning (ICML)*.

- [53] Honglei Zhang, Fangyuan Luo, Jun Wu, Xiangnan He, and Yidong Li. 2023. LightFR: Lightweight federated recommendation with privacy-preserving matrix factorization. *ACM Transactions on Information Systems* (2023), 1–28.
- [54] Jianqing Zhang, Yang Hua, Hao Wang, Tao Song, Zhengui Xue, Ruhui Ma, and Haibing Guan. 2023. FedALA: Adaptive Local Aggregation for Personalized Federated Learning. In *AAAI Conference on Artificial Intelligence (AAAI)*.
- [55] Jianqing Zhang, Dongjing Wang, and Dongjin Yu. 2021. TLSAN: Time-aware long-and short-term attention network for next-item recommendation. *Neurocomputing* 441 (2021), 179–191.
- [56] Michael Zhang, Karan Sapra, Sanja Fidler, Serena Yeung, and Jose M Alvarez. 2020. Personalized Federated Learning with First Order Model Optimization. In *International Conference on Learning Representations (ICLR)*.
- [57] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys* 52, 1 (2019), 1–38.
- [58] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-Level Convolutional Networks for Text Classification. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.
- [59] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep Leakage from Gradients. In *International Conference on Advances in Neural Information Processing Systems (NeurIPS)*.

A CONVERGENCE ANALYSIS

Recall that our objective is

$$\{\mathbf{W}_1, \dots, \mathbf{W}_N\} = \arg \min \mathcal{G}(\mathcal{F}_1, \dots, \mathcal{F}_N), \quad (9)$$

where $\mathcal{F}_i, \forall i \in [N]$ is the local loss and $\mathcal{G}(\mathcal{F}_1, \dots, \mathcal{F}_N) = \sum_{i=1}^N n_i \mathcal{F}_i$. During the training phase, the value of \mathcal{G} is the training loss of FedCP. To study the convergence of FedCP, we denote the loss calculated with the trained personalized models after local learning as $loss_{aft}$ and the loss calculated with the initialized personalized models before local learning as $loss_{bef}$. Except for the loss values, we also evaluate the corresponding test accuracy, calculated by averaging the accuracy of all the personalized models on the corresponding local test datasets of clients.

To empirically analyze the convergence of FedCP, we draw the training loss curves and test accuracy curves for our FedCP when using ResNet-18, as shown in Figure 6. On Tiny-ImageNet in the default practical setting, $loss_{aft}$ becomes close to $loss_{bef}$ after 74 iterations, and both of them reach the minimum value meanwhile. In other words, FedCP converges after training around 74 iterations. With the training loss decreasing, the test accuracy increases. Both the loss curve and the accuracy curve fluctuate before iteration 56 when using ResNet-18 due to the policy update, as shown in Figure 5 in the main body of this paper.




Figure 6: The training loss curves and test accuracy curves when using ResNet-18 on Tiny-ImageNet in the default practical setting. The red circles and green cubes represent the results evaluated before local learning and after local learning, respectively. Best viewed in color.

B PRIVACY-PRESERVING ABILITY

Here, following a traditional FL method FedCG [48], we consider a semi-honest scenario where the server follows the FL protocol but

may recover original data from a victim client with its model updates via Deep Leakage from Gradients (DLG) attack [59]. Among the baselines in our paper, there are two categories in terms of information transmission between the server and clients. Methods in Category 1 share the parameters in the entire backbone model, such as FedAvg, FedProx, Per-FedAvg, pFedMe, Ditto, FedRoD, FedFomo, and FedPHP. Methods in Category 2 only share the parameters in the feature extractor, such as FedPer and FedRep. Without loss of generality, we select the most famous methods in each category as the representative baselines: FedAvg for Category 1 and FedPer for Category 2. Also following FedCG, we provide the experimental results in Table 9 to evaluate the privacy-preserving ability of FedCP with representative baselines in Peak Signal-to-Noise Ratio (PSNR). The lower value of PSNR shows better privacy-preserving ability. The results in Table 9 show the superiority of FedCP.

Table 9: PSNR on Cifar100 in the default practical setting.

	FedAvg	FedPer	FedCP
PSNR (dB, ↓)	7.30	7.94	6.94

C CONDITIONAL POLICY NETWORK DESIGN

By default, our CPN consists of a fully connected (FC) layer [21] and a layer-normalization layer [3] (LN for short) followed by the ReLU activation function [28]. Here, we investigate how different designs affect the effectiveness of CPNs by varying the number of FC layers, the normalization layer, and the activation function, as shown in Table 10. Since the intermediate outputs $\mathbf{a}_i \in \mathbb{R}^{K \times 2}$ have two groups, we set the number of groups to two for the group-normalization [47] (GN for short). We only change the considered component based on FedCP. The accuracy results with an underline are higher than the accuracy of FedCP.

Table 10: The accuracy (%) with various CPNs on Tiny-ImageNet in the default practical setting.

	FedCP	Number of FC layers			Normalization layer		Activation function	
		/	2 FC	3 FC	4 FC	BN	GN	tanh
4-layer CNN	43.49	43.22	43.29	43.31	<u>44.13</u>	43.10	<u>43.89</u>	<u>43.92</u>
ResNet-18	44.18	<u>44.50</u>	<u>44.36</u>	44.11	43.70	43.25	43.49	<u>44.69</u>

The results in Table 10 show that we can further improve FedCP by using other architectures for the CPN. Adding more **FC layers** to process its input improves the test accuracy for ResNet-18 but causes a slight decrease for the 4-layer CNN. The additional parameters introduced for FedCP with 1 FC, 2 FC, 3 FC, and 4 FC are 0.527M (million), 0.790M, 1.052M, and 1.315M, respectively. However, the additional computing cost in each iteration introduced by additional FC layers is not worth the little accuracy increase. As for the **normalization layer**, replacing the LN with the batch-normalization [16] (BN) improves 0.64% test accuracy for the 4-layer CNN. However, it decreases around 0.48% accuracy for ResNet-18, which also contains BN layers. Similar to LN that normalizes entire \mathbf{a}_i , GN respectively normalizes $\mathbf{a}_{i,1}$ and $\mathbf{a}_{i,2}$. However, the test accuracy for both the 4-layer CNN and ResNet-18 decreases with the

GN layer. As for the activation function, using tanh only increases the accuracy for the 4-layer CNN, while using sigmoid improves the performance for both backbones compared to using ReLU, as the output belongs to (0, 1) is more suitable for outputting a policy.

D HYPERPARAMETER SETTINGS

We use the grid search to find the optimal λ . Specifically, we perform the grid search in the following search space:

- λ : 0, 0.1, 1, 5, 10

In this paper, we set $\lambda = 5$ for the 4-layer CNN and $\lambda = 1$ for the ResNet-18 and the fastText, respectively.

E DATA DISTRIBUTION VISUALIZATION

Here, we show visualizations of the data distributions (including training and test data) in the image and text tasks.

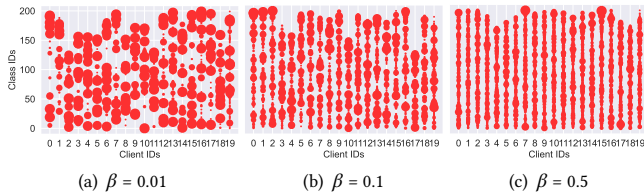


Figure 7: The data distribution of all clients on Tiny-ImageNet in practical settings with varying β . The size of a circle means the number of samples.

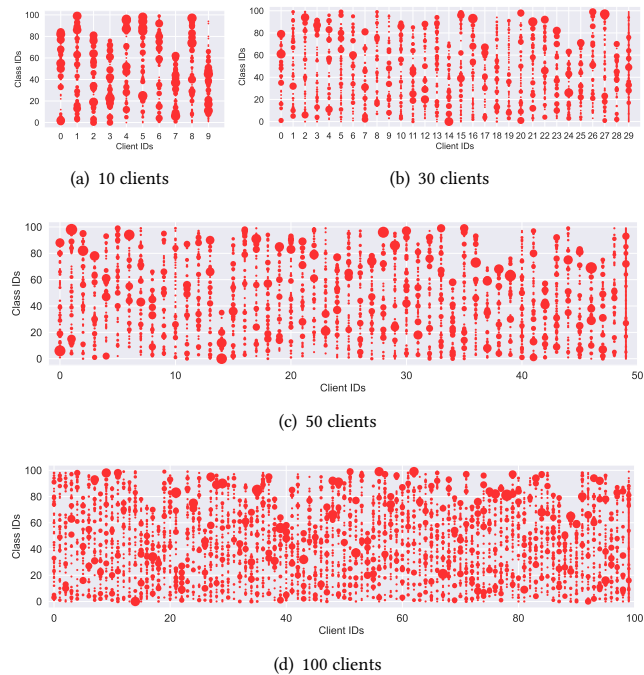


Figure 8: The data distribution of all clients on Cifar100 in practical settings with 10, 30, 50, and 100 clients, respectively.

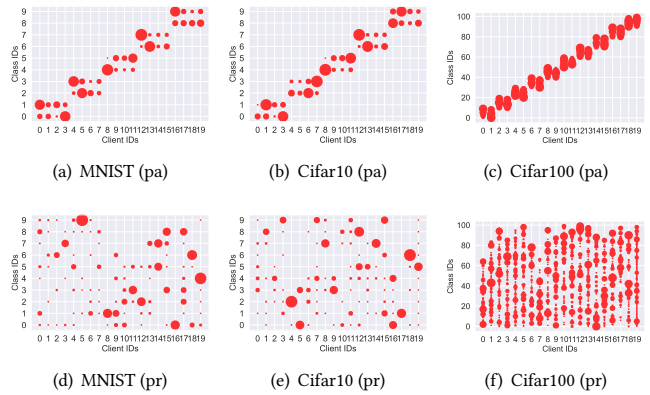


Figure 9: The data distribution of all clients in the pathological (pa) setting and default practical (pr) setting.

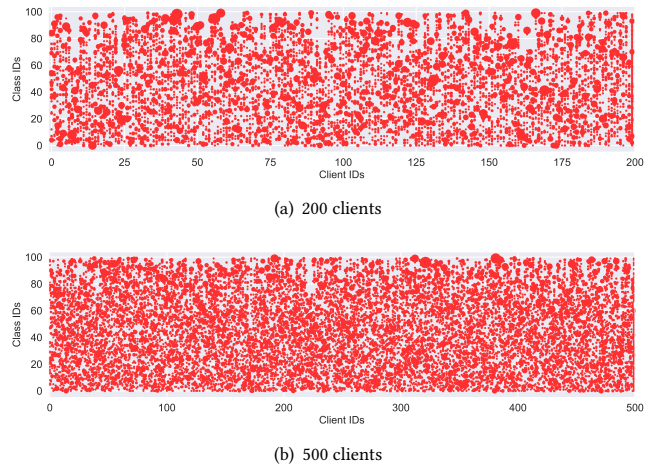


Figure 10: The data distribution of all clients on Cifar100 in default practical setting with 200 and 500 clients, respectively.




Figure 11: The data distribution of all clients on AG News in two heterogeneous settings.