

FedPall: Prototype-based Adversarial and Collaborative Learning for Federated Learning with Feature Drift

Yong Zhang^{1,3}, Feng Liang^{1*}, Guanghu Yuan¹, Min Yang², Chengming Li¹ and Xiping Hu^{1,3*}

Artificial Intelligence Research Institute, Shenzhen MSU-BIT University¹

Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences²

School of Medical Technology, Beijing Institute of Technology³

{zhangyong2023@bit.edu.cn, fliang@smbu.edu.cn, yuangh@mail.ustc.edu.cn
min.yang@siat.ac.cn, licm@smbu.edu.cn, huxp@smbu.edu.cn}

Abstract

Federated learning (FL) enables collaborative training of a global model in the centralized server with data from multiple parties while preserving privacy. However, data heterogeneity can significantly degrade the performance of the global model when each party uses datasets from different sources to train a local model, thereby affecting personalized local models. Among various cases of data heterogeneity, feature drift, feature space difference among parties, is prevalent in real-life data but remains largely unexplored. Feature drift can distract feature extraction learning in clients and thus lead to poor feature extraction and classification performance. To tackle the problem of feature drift in FL, we propose FedPall, an FL framework that utilizes prototype-based adversarial learning to unify feature spaces and collaborative learning to reinforce class information within the features. Moreover, FedPall leverages mixed features generated from global prototypes and local features to enhance the global classifier with classification-relevant information from a global perspective. Evaluation results on three representative feature-drifted datasets demonstrate FedPall’s consistently superior performance in classification with feature-drifted data in the FL scenario.¹

1. Introduction

Today, in computer vision, researchers often utilize large amounts of data from various parties to improve the accuracy of algorithms. However, this raises concerns, such as the potential for user privacy leakage caused by sharing private data [12]. Federated learning (FL) [22] is proposed as a

*Corresponding authors.

¹The code is available at <https://github.com/DistriAI/FedPall>.

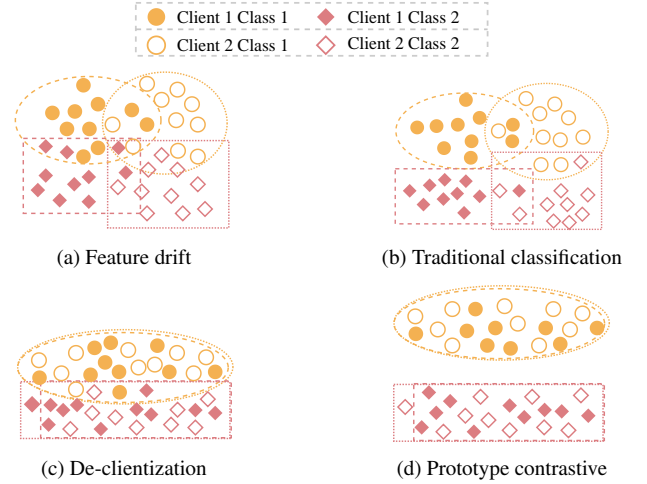


Figure 1. We show a schematic diagram of feature drift and use different techniques to drive feature distribution to update in different directions.

privacy-preserving distributed learning paradigm to address these challenges. In the FL paradigm, each party maintains a local client model and collaborates with others to train a global model on the server without sharing the original data, effectively protecting user privacy.

Particularly, the FL paradigm faces challenges from data heterogeneity [14]. Due to the limited local view of each client, data distribution discrepancies arise across clients (commonly referred to as the non-independent and identically distributed data issue, or non-independent and identically distributed (non-IID) data issue), which can increase generalization error for local models and degrade the performance of the globally aggregated model [10]. Although recent work on non-IID data in FL primarily addresses issues such as stability, client drift, and heterogeneous label distributions among clients [11, 18, 31], feature drift (i.e.,

variations in feature distributions across clients) is a prevalent yet underexplored challenge in FL. As shown in Fig. 1a, feature drift refers to the phenomenon where samples of the same class exhibit different feature distributions across different clients due to variations in data collection methods, devices, and other factors. This leads to ambiguous decision boundaries, severely impacting the classification performance of federated learning models. However, traditional classification losses (e.g., cross-entropy loss (CE loss)), as shown in Fig. 1b) do not account for feature drift. As a result, the collaborative effect among different clients causes the feature space of the same-class samples to influence each other, leading to only slight or even no clearer distinction between decision boundaries for various classes within the same client. As illustrated in Fig. 1c, decentralization is a common approach to addressing feature drift, which involves ignoring inter-client differences to make the feature space of same-class samples across different clients more clustered. In feature drift scenarios, some state-of-the-art algorithms do not explore the role of loss functions but instead optimize local update algorithms to achieve decentralization. For example, FedBN [19] compresses the feature space across clients by adding batch normalization layers to local models. Some methods [1, 34] align feature spaces by sharing partial data for synthetic data augmentation. Additionally, ADCOL sends raw features to the server to update an amplifier and uses the Kullback-Leibler (KL) loss function to achieve decentralization. However, these methods have drawbacks: on the one hand, they may lead to loss of class information, and on the other hand, they pose potential risks of privacy leakage.

To tackle the above challenge, we propose *FedPall*, a novel prototype-based adversarial and collaborative learning framework for FL with feature drift. FedPall applies adversarial learning between clients and the server to unify feature spaces, as well as collaborative learning across clients to enhance global decision boundaries. Specifically, FedPall uses adversarial learning to train a feature enhancer and uses KL divergence to align heterogeneous feature spaces between clients, while using prototype contrastive loss to reinforce class information (see Fig. 1d). Finally, adversarially aligned features are securely mixed with the global prototypes and uploaded to the server, where a global-view classifier is trained to enhance overall performance. Our contributions are summarized as follows:

- We propose a novel FL framework, FedPall, to address the feature drift problem. FedPall introduces adversarial learning between clients and the server, and collaborative learning among clients aiming to project feature representations into a unified feature space and reinforce the intrinsic class information. This approach effectively mitigates the feature drift problem in FL settings.
- We develop a technical strategy that hierarchically in-

tegrates the global prototypes with local features to orchestrate client-server collaboration. The mixed prototype features are then used to train a global classifier, which induces the classifier to distill discriminative patterns through cross-client knowledge consolidation.

- Empirical evaluation on three typical feature-drifted benchmarks demonstrates that our proposed method achieves state-of-the-art classification accuracy.

2. Related Work

In federated learning (FL), each client’s limited local view exacerbates the feature drift problem. Variations in data distributions can cause the same class label to have different feature representations, leading to poor generalization of local models. Existing solutions typically focus on two main approaches: feature space alignment and contrastive prototype learning.

Some studies address feature drift through feature representation alignment. FRAug [1] employs data augmentation to generate synthetic embeddings that capture both global and client-specific information. FedSea [26] aligns feature distributions to transform raw features into an IID format, mitigating feature drift. FedCiR [20] maximizes mutual information between representations and labels, while minimizing mutual information between client-specific representations conditioned on labels. FedHEAL [2] builds a fair aggregation target on the server side, indirectly aligning feature spaces. MOON [16] encourages local models to align with the global feature distribution through constrained updates. ADCOL [17], unlike traditional aggregation frameworks, uses adversarial learning to enforce a unified representation across clients, alleviating feature drift. However, ADCOL’s weak collaboration does not address global class boundaries, and its adversarial mechanism introduces potential privacy risks. Our approach enhances global class boundaries through stronger collaborative learning, while preserving privacy via adversarial learning to mitigate inter-client feature drift.

Some studies have explored prototype-driven federated learning (FL) paradigms, where prototypes compact feature embeddings, reducing communication bandwidth and preserving privacy [30]. Tan et al. [27, 28] proposed a supervised contrastive loss function using both global and local prototypes to minimize the distance between feature representations and class prototypes, addressing data heterogeneity. However, using the average feature as a class prototype may overlook intra-class variability. To address this, MP-FedCL [25] employs clustering on the client side to generate multiple prototypes per class, capturing intra-class variation and mitigating feature drift. Building on MP-FedCL, FedPLVM [29] improves local training with a two-stage clustering process between clients and the server, and incorporates an α -sparsity prototype loss to further op-

timize performance. By leveraging prototypes' privacy-preserving nature, this method addresses privacy concerns while reinforcing class-specific information in feature representations. In the FedPall framework, we enhance client-server collaboration through global prototypes. The server provides global category information to the client, helping bring similar data closer and separate dissimilar data. Additionally, we use mixed features with global category information to strengthen the global classifier.

3. Method

3.1. Problem Description

We define *feature drift* as follows: Given a dataset \mathbb{D} with features x and labels y , while the conditional distribution $P_i(x|y)$ differs across clients, the marginal distribution $P(y)$ remains the same. This means that the same label may have significantly different features across clients. For example, due to variations in environment, geographic location, and cultural differences, the structural features of houses can vary widely.

With feature drift in FL, our goal is to optimize each client's personalized model loss while leveraging the potential performance gains from collaborative learning across clients [33]. In the FedPall framework, there are a total of N clients, each client n has a private dataset D_n . Based on this goal, we formulate the overall optimization objective of the FedPall framework as follows:

$$\min_{\theta_1, \theta_2, \dots, \theta_N \in \mathbb{R}^{d_1}} F(\theta) := \frac{1}{N} \sum_{n=1}^N f_n(\theta_n), \quad (1)$$

where f_n represents the expected loss obtained from client n using the global model parameters under the dataset D_n , and θ_i represents the local model parameters of client i .

3.2. FedPall Framework

Existing approaches to addressing the feature drift problem in FL typically focus on either collaborative learning or adversarial learning in isolation. This can result in models that either fail to adequately capture class-related information in the feature representations or exhibit persistent discrepancies in feature distributions across clients. To address these limitations, we propose integrating both adversarial and collaborative learning to effectively mitigate feature drift in FL settings. In this section, we present the FedPall framework by elaborating on its adversarial and collaborative learning. The framework of the overall approach is shown in Fig. 2a. It is structured into four key procedures: generating global prototypes, training local models, training global models, and decentralizing global classifier.

We define certain model symbols here that will be used in this section. The local model $F(\cdot)$ consists of two components, a feature extractor $G(\cdot)$ (e.g., Resnet50 [7] for image

data) and a classifier $H(\cdot)$. We use a multilayer perceptron (MLP) as our Amplifier, and except for the output layer, the number of nodes in other layers is consistent with that of the classifier.

3.2.1. Generating Global Prototypes

Several studies [23, 27] suggest that category-centered prototypes are a privacy-friendly form of global knowledge. We leverage collaboration between clients to aggregate and generate global prototypes. Typically, each class prototype is defined as the mean of the feature vectors belonging to that class. The distinction between local and global prototypes lies in the scope of their underlying data. This section will describe the construction of both local and global prototypes.

The local prototype for the k -th category on client n is defined as:

$$c_n^k = \frac{1}{N_n^k} \sum_{(x,y) \in D_n^k} G_n(x), \quad (2)$$

where D_n^k and N_n^k represent the data samples and the number of samples for the k -th category on client n , respectively.

Gathering all the local prototypes together forms a local prototype set, which can be defined as:

$$\mathcal{O}_n = \{c_n^1, c_n^2, \dots, c_n^K\} \in \mathbb{R}^{K \times d}, \quad (3)$$

where K represents the number of categories owned by each client, and d denotes the output feature dimension.

Since we are only addressing the problem of feature drift, all clients have the same number of categories. Upon receiving the local prototype set and local label proportions $\mathcal{N} = \{\{N_n^k\}_{k=1}^K | n \in \mathcal{A}\}$ sent by client set \mathcal{A} , the server integrates the local prototypes from all clients to form the global prototypes,

$$\mathcal{G}^k = \sum_n \frac{N_n^k}{\sum_n N_n^k} c_n^k \quad (4)$$

The global prototypes set can be represented as

$$\mathcal{G} = [\mathcal{G}^1, \dots, \mathcal{G}^k, \dots, \mathcal{G}^K] \quad (5)$$

Next, the server sends the global prototype set \mathcal{G} to each client to guide local model training. You can refer to Fig. 2b for a better understanding of the generation of the global prototype.

3.2.2. Training Local Models

The goal of this module is to train an effective feature encoder that maps the raw data from different clients into a unified feature space, where the feature distributions are aligned and the class-related information is enhanced.

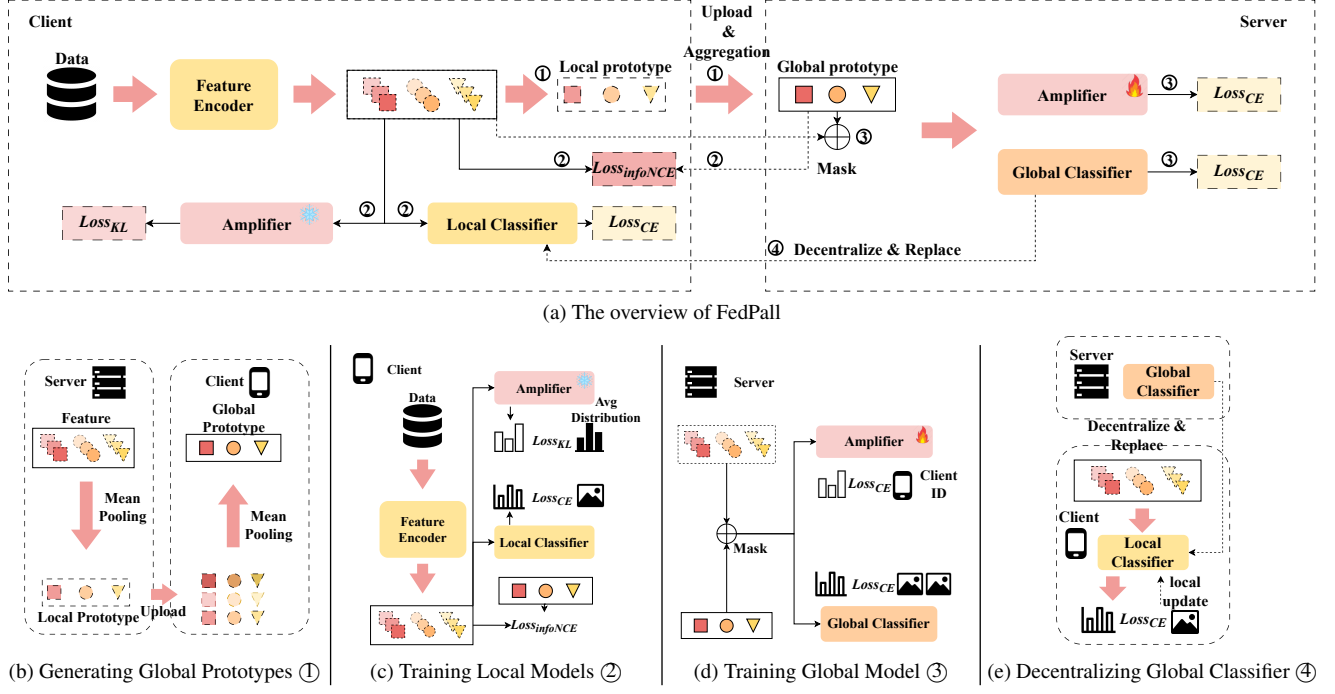


Figure 2. The FedPall framework and detailed phases

As mentioned earlier, due to feature drift, training a local classifier alone is not sufficient for accurately classifying data with feature drift. To address this, we apply adversarial learning to train a feature encoder. Specifically, we use a global amplifier, trained on the server, which amplifies the heterogeneous information in the features from different clients. At the client’s side, we apply the amplifier and use Kullback-Leibler (KL) divergence to reduce the heterogeneous information in the features, thereby creating an adversarial learning setup between the client and server. The objective is to improve the generalization ability of the feature encoders across clients while minimizing the client-specific heterogeneity in the feature representations. Let $x^{(i)}$ denote the i -th dimension of vector x , the KL divergence is calculated as:

$$\begin{aligned}\mathcal{L}_{KL} &= \sum_{(\mathbf{x}, \mathbf{y}) \in D_n} D_{KL}(A_n(G(\mathbf{x})) || [\frac{1}{N}]^N) \\ &= \sum_{(\mathbf{x}, \mathbf{y}) \in D_n} \sum_i A_n(G(\mathbf{x}))^{(i)} \log N A_n(G(\mathbf{x}))^{(i)}.\end{aligned}\quad (6)$$

where, $x^{(i)}$ denotes the i -th dimension of vector x , $A_n(\cdot)$ represents the amplifier of the n -th client, $D_{KL}(P||Q)$ represents the KL divergence of P and Q .

After adversarial learning, although the feature representations of different clients are mapped to similar feature distributions, the class-related information may be blurred. To address this problem, we propose using contrastive learning

to reinforce the class-related information within the feature encoder. To prevent information leakage, we enable collaboration between the global prototypes (server-side) and local features (client-side). Specifically, we employ the InfoNCE loss to minimize the distance between local features and their corresponding global prototypes, while maximizing the distance between local features and global prototypes of other classes. This approach strengthens the class-related representation within the feature encoder. The formula for the global-prototype contrastive loss is as follows:

$$\mathcal{L}_{infoNCE} = \sum_{(\mathbf{x}, \mathbf{y}) \in D_n} -\log \frac{\exp(\text{sim}(G(\mathbf{x}), \mathcal{G}^y)/\tau)}{\sum_{y_\alpha \in A(y)} \exp(\text{sim}(\mathbf{x}, \mathcal{G}^{(y_\alpha)})/\tau)}, \quad (7)$$

where $A(y) := \{y_\alpha \in [1, |\mathcal{G}|] : y_\alpha \neq y\}$ is the set of labels distinct from y , τ is the temperature to adjust the tolerance for feature difference, and $\text{sim}(x, y)$ represents the cosine similarity of x and y .

We combine adversarial learning and collaborative learning to address the feature drift problem in FL. By leveraging the two loss functions defined above, along with the local cross-entropy loss, we progressively train the local feature encoder at each client. The overall loss function for each client is as follows:

$$\mathcal{L} = \mathcal{L}_{CE(x, y) \sim D_n}(F(\mathbf{x}), y) + \mu \mathcal{L}_{KL} + \delta \mathcal{L}_{infoNCE}, \quad (8)$$

where μ denotes the weight of the \mathcal{L}_{KL} divergence and \mathcal{L}_{CE} denotes the cross-entropy loss.

The local model $F(\cdot)$ is updated using Eq. (8). For a comprehensive visualization of the global prototype generation process, consult Fig. 2c where the workflow is systematically delineated.

3.2.3. Training Global Model

Due to the client's limited local view, training an accurate classifier is challenging. To address this, encrypted mixed features are uploaded to the server, enabling classifier training with a global perspective. Additionally, the server leverages its global view to train an amplifier, which engages in adversarial learning with the client to enhance model robustness. The detailed process is described below.

For each feature z_n^k of class k on client n , we obtain a prototype mixed feature by performing a weighted fusion with the corresponding global prototype:

$$r_n^k = \alpha \times z_n^k + (1 - \alpha) \times \mathcal{G}^k, \quad (9)$$

where $z_n^k \in Z_n^k = \{z_n^{i,k}\}_{i=1}^S$. $\alpha \sim U(u_f, u_r)$ are the mix parameters, with u_f and u_d representing two hyperparameters of a uniform distribution. S represents the number of samples of the k -th category in client n .

Building on this, we employ a Bernoulli mask to further reduce the risk of privacy leakage,

$$\begin{aligned} \text{Mask} &= \{X_1, X_2, \dots, X_d\}, \\ X_i &\sim \text{Bernoulli}(\beta), \quad \forall i \in [1, d], \end{aligned} \quad (10)$$

The final output of the prototype mixing mechanism is derived by selecting the mixed feature elements based on the noise mask:

$$\tilde{r}_n^k = \text{Mask} \odot r_n^k, \quad (11)$$

where \odot is an element-wise *and* operator.

After generating the prototype mixed features, the client will form the set $\mathcal{D}_{R_L}(R, Y)$ using the prototype mixed feature set $R_n = \{\tilde{r}_n^1, \dots, \tilde{r}_n^k, \dots, \tilde{r}_n^K\} \in \mathbb{R}^{K \times d}$ and the corresponding label set Y , which will be sent to the server.

The server updates the global amplifier A using the mixed feature sets from each client along with the corresponding client IDs. Specifically, we first construct the dataset for training the amplifier, denoted as $\mathcal{D}_{R_I}(R, I)$, where I represents the client IDs. The amplifier is then updated by minimizing the empirical risk:

$$\mathbb{E}_{(R, I) \sim \mathcal{D}_{R_I}} \ell_{CE}(R, I). \quad (12)$$

At the same time, the server updates the global classifier C_g using the mixed prototype features and the class labels from the clients Y . This is done by minimizing the empirical risk:

$$\mathbb{E}_{(R, Y) \sim \mathcal{D}_{R_L}} \ell_{CE}(R, Y). \quad (13)$$

We show the training process of the global classifier in Fig. 1c.

3.2.4. Decentralizing Global Classifier

Finally, we deploy the global classifier C_g to each client to replace the original local classifiers C_c . The purpose of this is to obtain a more generalized classifier that can alleviate the feature drift problem. To allow the global classifier to adapt to the personalized characteristics of local data, we retrain it on the client's local data, thereby enhancing the classifier's accuracy and improving its performance on individual client data distributions. And you can understand the deployment method of the server-side global classifier through Fig. 2e.

4. Discussion

Computational Cost Compared to the standard federated learning model, our adversarial collaborative learning approach introduces an amplifier and a global classifier. However, the number of parameters for these two components is much smaller than those of the other components. In our design, both the amplifier and the classifier are designed as three-layer MLPs. Compared to the feature extractor (with a total of 23.5M parameters), the classifier (with a total of 1.32M parameters) and the amplifier (with a total of 1.33M parameters) account for approximately 5.61% and 5.59%, respectively. Moreover, the client-side amplifier remains frozen, functioning exclusively in the forward pass for loss calculation while being disabled during backpropagation cycles.

Communication Efficiency Unlike traditional federated learning methods such as FedAvg, which transmit full local model parameters, our approach uses prototype-mixed features as the communication medium. Assuming consistent tensor representations between model parameters and prototype features, each client employs a ResNet-50-based feature extractor that generates 2048-dimensional embeddings. During upload, clients transmit an average of 12,122 prototype-mixed features, compared to approximately 24.8 million model parameters in FedAvg. For downloading, only the amplifier and global classifier parameters are transmitted, reducing communication cost to about 10.6% of that in FedAvg (see previous paragraph).

Limitation While our framework currently specializes in image recognition tasks, its extension to NLP or time-series analysis remains unexplored. Successful cross-domain adaptation requires two key developments: (1) establishing domain-specific feature representations and prototype definitions, and (2) redesigning loss functions according to task semantics. For NLP applications, this implies reconfiguring the standard classification paradigm into autoregressive prediction frameworks. Architectural adaptations are equally crucial - particularly the incorporation of RNN-based struc-

tures with inherent temporal modeling capabilities for sequential data processing.

5. Experiments

5.1. Experimental Setup

Datasets We conduct experiments on three publicly available feature drift datasets: Digits [32], Office-10 [6], and PACS [15]. Specifically, (1) the Digits dataset consists of five different domain sources: MNIST [13], SVHN [24], USPS [9], SynthDigits [5], and MNIST-M [5]; (2) the Office-10 dataset includes four distinct sources: Amazon, Caltech, DSLR, and WebCam; (3) the PACS dataset consists of four sources: Art Painting, Cartoon, Photo, and Sketch. Datasets Office-10 and PACS are real-world images from natural scenes, which inherently exhibit feature drift due to the diversity of their sources. Digits is a digit recognition dataset. In line with [19, 28], we do not use the entire Digits dataset for feature transformation experiments but rather a subset of 10% of the data. For datasets Office-10 and PACS, we used all of the datasets for the experiment. Additionally, we split each dataset into training and testing sets with an 8:2 ratio.

Baselines We compare FedPall with ten baselines, including SingSet (where each client independently trains a model). FedAvg [22] is the most classic federated learning algorithm, while FedProx [18], PerFedAvg [4], and FedRep [3] are personalized federated learning methods. FedBN [19], ADCOL [17], MOON [16], FedProto [27] and FedHEAL [2] are personalized federated learning algorithms for cross-domain learning, all of which address the issue of non-IID features to some extent. In addition, we explored the ability of RUCR [8] to solve the feature drift problem.

Model and Hyper-parameter Setup All algorithms adopt identical local model architectures for fair comparison. Each local model contains: (1) a ResNet-50 feature extractor (excluding classifier layer), (2) a three-layer MLP classifier with 512 hidden units, and (3) a three-layer MLP amplifier with 2048 input/512 hidden units. Output dimensions for classifier and amplifier are set according to dataset categories and data sources, respectively. We maintain client count equal to data sources (one source per client). Local training uses 5 epochs for Digits, and 10 epochs for Office-10 and PACS. Global training employs 100 epochs throughout. Optimization uses SGD ($\text{lr}=0.01$). Except for the digits dataset, for which the values of μ and δ are set to 0.7 and 0.3 respectively, the values of μ and δ are set to 0.1 for all other datasets. Common loss hyperparameters remain fixed across algorithms. All experiments run on NVIDIA RTX 4090 GPUs.

5.2. Experimental Analysis

We conduct the evaluation on three publicly available feature-drifted datasets (Digits, Office-10, and PACS) and compare the performance of the FedPall framework with classical and state-of-the-art baselines. As shown in Table 1, our proposed framework achieves state-of-the-art accuracy on all three datasets.

We first discuss the experimental results based on each individual dataset. On the Office-10 dataset, the overall accuracy of the FedPall framework surpasses that of the second-best method, ADCOL, by approximately 3 percentage points. On the Digits dataset, it is evident that FedPall outperforms all other models, achieving an accuracy that is approximately 1.1 percentage points higher than the second-best model, FedBN. The Digits dataset contains images that are relatively easy to classify, and the degree of feature drift is smaller compared to the Office-10 dataset. All baseline models achieve reasonably good accuracy on this dataset. Specifically, adversarial learning helps mitigate the heterogeneous information in the MNIST-M client. Similarly, our algorithm demonstrates strong performance on the PACS dataset, achieving an overall accuracy that is approximately 1.1 percentage points higher than the second-highest result produced by FedBN. FedPall achieves the highest or second-highest accuracy across all sub-datasets.

We also discuss the performance of FedPall as compared to other state-of-the-art baselines across all three datasets. The average accuracy of FedPall consistently outperforms that of ADCOL in all three datasets, achieving an increase ranging from about 1.1 to 2.9 percentage points. In addition, even though FedBN can achieve accuracy comparable to our method on datasets Digits and PACS, our method outperforms it significantly by 31.5 percentage points on datasets Office-10. As mentioned earlier, the Office-10 dataset comes from real-world data, where feature drift is particularly prominent, and there is also a significant distribution difference between the training and testing sets, leading to the suboptimal performance of the FedBN method on this dataset. In contrast, the special design incorporating both adversarial and collaborative learning in FedPall enables it to adapt well to the Office-10 dataset.

5.3. Ablation Study

Effect of loss combination In this section, we analyze the impact of KL loss and InfoNCE loss on the performance of the local feature encoder. We conduct ablation studies with three configurations: (1) removing both KL and InfoNCE losses, (2) removing only KL loss, and (3) removing only InfoNCE loss. As shown in Fig 3, the algorithm performs best when all losses are retained, which validates the reliability of the loss combination we designed.

Specifically, on the Office-10 dataset, our method out-

Table 1. The top-1 accuracy (%) of each algorithm on each sub-dataset of the Office-10, Digits, and PACS datasets is compared, along with the average top-1 accuracy across all sub-datasets. The mean and standard deviation (std) from three random trials (using different random seeds, with other experimental settings remaining the same) are reported. The highest accuracy for each dataset is highlighted in bold, and the second-highest accuracy is underlined.

		SingleSet	FedAvg	FedProx	PerfedAvg	FedRep	FedBN	MOON	FedProto	ADCOL	RUCR	FedHEAL	ours(FedPal)
Office-10	amazon	74.0(2.7)	56.9(2.5)	56.6(2.6)	57.1(2.2)	45.3(1.9)	40.8(15.8)	51.7(16.1)	69.4(2.1)	73.3(4.4)	52.1(8.5)	65.1(3.3)	72.9(1.4)
	caltech	44.7(3.2)	46.5(4.6)	51.0(5.2)	50.8(1.6)	38.4(4.9)	33.9(6.5)	41.3(13.6)	39.4(6.3)	37.2(1.7)	44.3(1.0)	44.6(3.2)	44.7(8.7)
	dslr	60.2(6.7)	30.1(4.9)	33.3(10.4)	31.2(4.9)	34.4(4.9)	38.7(3.2)	24.7(1.9)	65.6(4.9)	76.3(4.9)	30.1(6.7)	67.7(1.9)	77.4(3.2)
	webcam	71.3(2.6)	37.9(6.2)	43.7(7.2)	47.1(7.8)	55.8(2.6)	30.5(6.1)	33.3(12.7)	71.3(4.3)	71.3(2.6)	37.4(5.0)	60.9(1.0)	74.7(1.0)
	avg	62.5(0.4)	42.9(1.2)	46.1(2.6)	46.6(2.9)	43.5(1.3)	36.0(6.5)	37.8(10.9)	61.4(1.7)	64.5(1.8)	41.0(0.6)	59.6(0.5)	67.5(2.7)
Digits	MNIST	95.5(0.2)	92.9(2.2)	91.8(3.0)	90.1(4.8)	86.5(6.1)	96.7(0.1)	93.4(1.1)	96.4(0.5)	96.3(0.4)	92.6(2.0)	93.6(0.6)	97.2(0.4)
	SVHN	71.1(0.9)	77.4(0.2)	76.9(0.3)	75.6(0.4)	67.2(1.7)	79.4(0.3)	79.6(0.8)	72.5(0.3)	75.1(2.1)	77.9(0.3)	68.3(2.0)	78.0(0.4)
	USPS	86.4(0.3)	89.3(0.9)	89.2(1.4)	88.7(0.7)	90.0(3.0)	90.1(0.5)	81.8(0.7)	87.0(0.8)	86.7(1.3)	88.9(2.4)	87.0(0.5)	87.3(1.3)
	SynthDigits	95.2(0.1)	95.5(0.1)	95.4(0.1)	95.0(0.2)	94.2(0.8)	95.6(0.1)	96.6(0.2)	95.3(0.6)	96.4(0.3)	96.0(0.2)	89.3(1.6)	95.3(0.4)
	MNIST-M	76.6(0.4)	73.8(1.5)	74.0(1.5)	73.2(0.8)	69.1(0.9)	76.3(0.4)	72.2(0.9)	78.3(1.2)	78.3(4.4)	72.7(0.4)	67.8(1.9)	85.9(1.4)
	avg	84.9(0.1)	85.8(0.9)	85.5(1.1)	84.5(1.3)	81.4(2.5)	87.6(0.1)	84.7(0.6)	85.9(0.2)	86.6(1.3)	85.6(0.9)	81.2(1.3)	88.7(0.2)
PACS	art.painting	33.6(0.8)	25.8(1.9)	24.3(4.1)	26.5(2.2)	26.9(3.3)	36.7(1.8)	30.6(2.0)	32.7(0.7)	34.9(1.2)	24.7(1.1)	31.2(1.2)	35.6(0.6)
	cartoon	58.5(2.5)	45.4(2.3)	51.4(0.6)	48.3(1.2)	44.4(2.1)	55.6(2.0)	51.5(1.8)	57.3(1.5)	57.2(0.8)	47.5(3.3)	50.8(0.4)	59.7(2.3)
	photo	63.0(1.9)	48.7(3.1)	49.6(2.0)	46.9(2.6)	41.9(2.8)	66.1(1.0)	53.0(3.3)	64.0(1.3)	62.1(2.0)	47.5(6.2)	61.1(2.0)	64.7(1.3)
	sketch	79.7(0.1)	49.0(2.0)	40.7(1.5)	44.4(3.8)	40.5(1.3)	79.6(1.7)	55.1(1.4)	79.6(0.8)	80.1(1.0)	42.2(2.4)	73.8(0.1)	82.2(0.7)
	avg	58.7(1.2)	42.2(1.6)	41.5(1.8)	41.5(1.9)	38.4(1.1)	59.5(1.4)	47.6(0.9)	58.4(0.3)	58.6(0.6)	40.5(2.0)	54.2(0.5)	60.6(0.4)




Figure 3. We evaluate the top-1 accuracy averaged over all clients using different loss function combinations on different datasets.

performs the model trained with only CE loss by nearly 4 percentage points. Notably, combining CE loss with InfoNCE loss yields worse results than using CE loss alone, suggesting that in the presence of severe feature drift, reinforcing category information through InfoNCE may amplify the drift. While the CE + KL loss combination performs better than CE loss alone on Office-10, it underperforms on the PACS dataset—even falling below CE loss—indicating reduced robustness. This suggests that CE + KL may compromise category information, leading to instability across heterogeneous datasets.

To assess the impact of KL and InfoNCE losses on feature distributions, we visualize class-wise features across clients in the Office-10 dataset using t-SNE projections of randomly sampled data points (Fig 4). The CE-only model poorly mitigates feature drift, with Client 1 showing unclear decision boundaries. Adding InfoNCE improves intra-client class separation but fails to resolve inter-client drift, leading to ambiguous global boundaries. The CE+KL combination reduces cross-client distances for the same class, yielding clearer global boundaries; however, it compresses

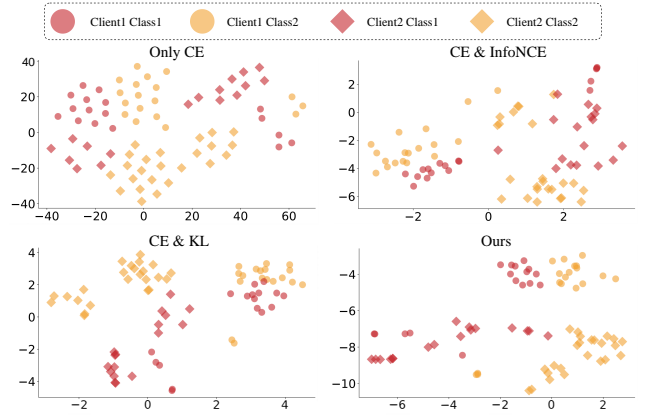


Figure 4. We plotted the feature distribution of different categories under different clients, corresponding to the four loss combination strategies of Fig. 3

intra-class spacing in Client 1, causing overlapping clusters and outliers that hinder local classification. In contrast, our unified loss balances these effects: KL aligns same-class features across clients, while CE and InfoNCE promote intra-client separation. This coordination produces compact, well-separated clusters, improving classification and validating our method’s effectiveness.

In summary, our method outperforms other loss combinations in the simple handwritten digit recognition task and achieves superior results on real-world datasets, demonstrating strong robustness and generalization.

Hyperparameter sensitivity analysis. We conducted a hyperparameter sensitivity analysis on the Office-10 dataset by varying the loss weights μ and δ over $\{0.1, 0.2, 0.5, 1.0\}$. To visualize their impact, we generated an interpolated heatmap of average accuracy. As shown in Fig 5, ac-




Figure 5. The top-1 accuracy (%) of different μ and δ under Office-10 dataset

accuracy remains high and stable when μ is within $[0.1, 0.4]$, with δ having minimal influence in this range. Notably, the highest accuracy of 69.12% occurs at $\mu = \delta = 0.2$, followed by 68.62% at $\mu = 1.0, \delta = 0.2$.

The results show that our algorithm is not sensitive to the parameter changes of the loss function within a certain range and can maintain high performance.




Figure 6. Comparison of accuracy with and without training the global classifier on the three datasets.

Comparison of different classifier replacement methods We evaluate the effectiveness of the global classifier through an ablation study by removing it. As shown in Fig 6, although the simplicity of the Digits dataset results in slightly lower accuracy for the global classifier on some sub-datasets (panel 6a), it still outperforms the baseline without a global classifier across multiple sub-datasets. Notably, our method achieves 3.6 percentage

points higher accuracy on MNIST-M. The benefits are more pronounced on datasets with substantial feature drift: on Office-10 (panel 6b), the global classifier surpasses the baseline on nearly all sub-datasets, achieving 49.3% on Caltech—a 15.6 percentage points improvement. On PACS (panel 6c), it consistently outperforms the baseline across all sub-datasets, with gains of up to 3 percentage points.

These results confirm the necessity of FedPall’s global classifier, which captures cross-client category information to enhance client-server collaboration and improve the framework’s generalization against feature drift.

Privacy Leakage Risk We evaluate the privacy risk of prototype mixture features using the Data Efficient Mutual Information Neural Estimator (DEMINE) [21]. On the Office-10 training set, the mutual information (MI) scores for: (1) standard Gaussian noise, (2) prototype mixture only, (3) Bernoulli masking only, and (4) our method (prototype mixture + Bernoulli masking) are 2.96, 3.06, 3.04, and **2.10**, respectively (lower is better). The corresponding average accuracies are 66.4%, 64.5%, 63.7%, and **67.6%**, compared to 65.8% in the noise-free case. Our approach not only offers stronger privacy protection but also improves accuracy, attributed to the consistent update direction between the global prototype and feature encoder. Additionally, since encryption is applied post-training with linear time complexity, the computational overhead remains negligible.

6. Conclusion

In this study, we focus on the feature drift problem in FL. The feature drift problem causes the same class samples on different clients to have distinct feature distributions, making it difficult for traditional model aggregation methods to handle such data heterogeneity. To tackle this problem, we design a prototype-based adversarial collaborative framework to unify feature spaces and enhance classification boundaries. The global classifier is retrained with mixed features to further grasp classification-relevant information from a global perspective. Our method has empirically achieved state-of-the-art performance in popular feature-drifted datasets with multiple data sources.

Acknowledgement

This work was supported in part by the Innovation Team Project of Guangdong Province of China (2024KCXTD017) and Guangdong-Hong Kong-Macau University Joint Laboratory (2023LSYS005).

References

- [1] Haokun Chen, Ahmed Frikha, Denis Krompass, Jindong Gu, and Volker Tresp. Fraug: Tackling federated learning with non-iid features via representation augmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4849–4859, 2023. 2
- [2] Yuhang Chen, Wenke Huang, and Mang Ye. Fair federated learning under domain skew with local consistency and domain diversity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12077–12086, 2024. 2, 6
- [3] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, pages 2089–2099. PMLR, 2021. 6
- [4] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in neural information processing systems*, 33:3557–3568, 2020. 6
- [5] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189. PMLR, 2015. 6
- [6] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2066–2073. IEEE, 2012. 6
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3
- [8] Wenke Huang, Yuxia Liu, Mang Ye, Jun Chen, and Bo Du. Federated learning with long-tailed data via representation unification and classifier rectification. *IEEE Transactions on Information Forensics and Security*, 2024. 6
- [9] Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994. 6
- [10] Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019. 1
- [11] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank J Reddi, Sebastian U Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for on-device federated learning. *arXiv preprint arXiv:1910.06378*, 2(6), 2019. 1
- [12] Jakub Konečný, Brendan McMahan, and Daniel Ramage. Federated optimization: Distributed optimization beyond the datacenter. *arXiv preprint arXiv:1511.03575*, 2015. 1
- [13] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 6
- [14] Boyuan Li, Shengbo Chen, and Zihao Peng. New generation federated learning. *Sensors*, 22(21):8475, 2022. 1
- [15] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 6
- [16] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10713–10722, 2021. 2, 6
- [17] Qinbin Li, Bingsheng He, and Dawn Song. Adversarial collaborative learning on non-iid features. In *International Conference on Machine Learning*, pages 19504–19526. PMLR, 2023. 2, 6
- [18] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020. 1, 6
- [19] Xiaoxiao Li, Meirui JIANG, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. In *International Conference on Learning Representations*. 2, 6
- [20] Zijian Li, Zehong Lin, Jiawei Shao, Yuyi Mao, and Jun Zhang. Fedcir: Client-invariant representation learning for federated non-iid features. *IEEE Transactions on Mobile Computing*, 2024. 2
- [21] Xiao Lin, Indranil Sur, Samuel A Nastase, Ajay Divakaran, Uri Hasson, and Mohamed R Amer. Data-efficient mutual information neural estimator. *arXiv preprint arXiv:1905.03319*, 2019. 8
- [22] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 1, 6
- [23] Xutong Mu, Yulong Shen, Ke Cheng, Xueli Geng, Jiaxuan Fu, Tao Zhang, and Zhiwei Zhang. Fedproc: Prototypical contrastive federated learning on non-iid data. *Future Generation Computer Systems*, 143:93–104, 2023. 3
- [24] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Baolin Wu, Andrew Y Ng, et al. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, page 4. Granada, 2011. 6
- [25] Yu Qiao, Md Shirajum Munir, Apurba Adhikary, Huy Q Le, Avi Deb Raha, Chaoning Zhang, and Choong Seon Hong. Mp-fedcl: Multi-prototype federated contrastive learning for edge intelligence. *IEEE Internet of Things journal*, 2023. 2
- [26] Min Tan, Yinfu Feng, Lingqiang Chu, Jingcheng Shi, Rong Xiao, Haihong Tang, and Jun Yu. Fedsea: Federated learning via selective feature alignment for non-iid multimodal data. *IEEE Transactions on Multimedia*, 2023. 2
- [27] Yue Tan, Guodong Long, Lu Liu, Tianyi Zhou, Qinghua Lu, Jing Jiang, and Chengqi Zhang. Fedproto: Federated prototype learning across heterogeneous clients. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8432–8440, 2022. 2, 3, 6
- [28] Yue Tan, Guodong Long, Jie Ma, Lu Liu, Tianyi Zhou, and Jing Jiang. Federated learning from pre-trained models: A contrastive learning approach. *Advances in neural information processing systems*, 35:19332–19344, 2022. 2, 6

- [29] Lei Wang, Jieming Bian, Letian Zhang, Chen Chen, and Jie Xu. Taming cross-domain representation variance in federated prototype learning with heterogeneous data domains. *arXiv preprint arXiv:2403.09048*, 2024. [2](#)
- [30] Li Wang, Quanguai Zhang, Lei Sang, Qiang Wu, and Min Xu. Federated prototype-based contrastive learning for privacy-preserving cross-domain recommendation. *arXiv preprint arXiv:2409.03294*, 2024. [2](#)
- [31] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. [1](#)
- [32] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI 16*, pages 561–578. Springer, 2020. [6](#)
- [33] Tailin Zhou, Jun Zhang, and Danny HK Tsang. Fedfa: Federated learning with feature anchors to align features and classifiers for heterogeneous data. *IEEE Transactions on Mobile Computing*, 2023. [3](#)
- [34] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International conference on machine learning*, pages 12878–12889. PMLR, 2021. [2](#)