

# Cooperative Pseudo Labeling for Unsupervised Federated Classification

Kuangpu Guo<sup>1,2</sup>, Lijun Sheng<sup>1,2</sup>, Yongcan Yu<sup>2,3</sup>, Jian Liang<sup>\* 2,3</sup>, Zilei Wang<sup>1</sup>, Ran He<sup>2,3</sup>

<sup>1</sup> University of Science and Technology of China

<sup>2</sup> NLPR & MAIS, Institute of Automation, Chinese Academy of Sciences

<sup>3</sup> University of Chinese Academy of Sciences

gkp@mail.ustc.edu.cn, liangjian92@gmail.com

## Abstract

*Unsupervised Federated Learning (UFL) aims to collaboratively train a global model across distributed clients without sharing data or accessing label information. Previous UFL works have predominantly focused on representation learning and clustering tasks. Recently, vision language models (e.g., CLIP) have gained significant attention for their powerful zero-shot prediction capabilities. Leveraging this advancement, classification problems that were previously infeasible under the UFL paradigm now present promising new opportunities, yet remain largely unexplored. In this paper, we extend UFL to the classification problem with CLIP for the first time and propose a novel method, **Federated Cooperative Pseudo Labeling (FedCoPL)**. Specifically, clients estimate and upload their pseudo label distribution, and the server adjusts and re-distributes them to avoid global imbalance among classes. Moreover, we introduce a partial prompt aggregation protocol for effective collaboration and personalization. In particular, visual prompts containing general image features are aggregated at the server, while text prompts encoding personalized knowledge are retained locally. Extensive experiments demonstrate the superior performance of our FedCoPL compared to baseline methods. Our code is available at <https://github.com/krumpguo/FedCoPL>.*

## 1. Introduction

Federated learning (FL) [24, 41] is a distributed machine learning paradigm that facilitates collaboration among clients while preserving data privacy by avoiding data sharing [23, 43]. Conventional FL predominantly follows a supervised learning paradigm [2, 8, 9, 53, 62], relying on well-annotated labels at the client side. However, data annotation is often time-consuming and requires domain expertise [32, 40, 42, 59], posing challenges in real-

world applications. Consequently, substantial research efforts have been dedicated to unsupervised federated learning (UFL) [13, 17, 48, 54], which facilitates collaborative model training among clients using unlabeled data. Existing UFL works [27, 37, 58] primarily focus on leveraging unlabeled data for representation learning or clustering, enabling models to extract features that enhance generalization and performance on downstream tasks. However, the absence of labeled data prevents the direct application of UFL to classification tasks.

Vision-language models (VLMs) [1, 22, 39], particularly contrastive language-image pre-training (CLIP) [39], have recently garnered significant attention for their remarkable representation and generalization capabilities [7, 33, 50]. By pre-training on large-scale image-text pairs, CLIP extracts modality-specific features through two encoders, enabling powerful zero-shot classification. The impressive zero-shot prediction capability of CLIP alleviates the dependency on labeled data. This advancement presents new opportunities for addressing previously intractable classification tasks in UFL, yet remains largely unexplored.

In this paper, we are the first to extend UFL to classification tasks by leveraging CLIP. Each client is initialized with the same pre-trained CLIP model and possesses only unlabeled local data. Similar to standard federated learning frameworks [51, 56], clients can periodically upload their non-data knowledge (e.g., statistical information, model parameters) to the central server for aggregation and receive updated global knowledge in return. However, when directly applied within the UFL framework, the inherent bias of CLIP and the unknown label distribution among clients [14, 32, 35] can result in globally imbalanced training across categories. Moreover, the potential label skew can introduce discrepancies in model update directions among clients, rendering conventional aggregation methods ineffective for collaboration among clients.

To address these challenges, we propose a novel method, **FedCoPL**, which comprises two key components: cooperative pseudo labeling and partial prompt aggregation. The

\*Corresponding author

cooperative pseudo labeling mechanism aims to construct a more representative pseudo label distribution while ensuring balanced training across all classes. At the beginning of federated training, each client leverages confidence-based and entropy-based filtering methods to estimate its pseudo label distribution, which is then transmitted to the server. The server globally adjusts the distribution for each client based on the proportion of pseudo labels in each category. This process alleviates class bias caused by CLIP’s inherent bias and the unknown data distribution among clients, ultimately enhancing the quality of pseudo labels.

To promote effective collaboration and personalization, we introduce a partial prompt aggregation strategy. Specifically, for efficient computation and communication, we only optimize the visual and textual prompts instead of the entire model parameters. Previous studies [52] have shown that visual prompts tend to capture general representations of the feature space, whereas textual prompts are more focused on category-specific information. Aggregating prompts from both modalities simultaneously may lead to parameter conflicts, especially in the presence of label skew. Therefore, we propose transmitting only visual prompts to the server for aggregation, while retaining textual prompts locally. This strategy enhances global collaboration through aggregated visual prompts while enabling clients to benefit from personalized textual prompts, which better align with their local label distributions. Experimental results on standard federated prompt learning benchmarks, with both Dirichlet-based and quantity-based label skews, demonstrate the effectiveness of the proposed FedCoPL compared to baseline methods. Our main contributions are summarized as follows:

- This paper extends the unsupervised federated learning paradigm to classification tasks with CLIP, which remains largely unexplored in previous works.
- We propose FedCoPL to mitigate label skew and the inherent bias of CLIP. FedCoPL primarily incorporates a cooperative pseudo labeling strategy to construct more representative and accurate pseudo labels.
- FedCoPL further integrates a partial prompt aggregation protocol to facilitate effective collaboration while preserving client personalization.
- Extensive experiments across six commonly used datasets with Dirichlet-based and quantity-based label skews demonstrate the effectiveness of FedCoPL.

## 2. Related Work

### 2.1. Unsupervised Federated Learning

Unsupervised federated learning aims to collaboratively train a global model across distributed clients with unlabeled data while preserving data privacy by avoiding data sharing. Due to the lack of prior label knowledge, exist-

ing unsupervised federated methods primarily concentrate on representation learning and clustering tasks. Specifically, both FedU [65] and FedEMA [66] enhance the awareness of heterogeneity in unsupervised federated learning by divergence-aware predictor update rule, and adaptive global knowledge interpolation, respectively. Additionally, FedUL [28] introduces federated learning with only unlabeled data but requires knowledge of precise label frequencies for each client. L-DAWA [40] proposes prototypical representation distillation to learn a representation from unlabeled data. FedUU [27] uses flexible uniform regularization to obtain uniform and unified representation and FedGaLA [37] performs gradient alignment at the client level to encourage clients to learn domain-invariant features. Moreover, Orchestra [30] utilizes local-global clustering to guide self-supervised learning. In contrast, we propose leveraging CLIP’s zero-shot capability to generate pseudo labels for model training, enabling the handling of more complex tasks like image classification.

### 2.2. Federated Learning for VLMs

The fine-tuning of vision-language models recently has been extended to the federated framework to alleviate the computational load while addressing challenges in federated learning, such as robustness in cross-domain scenarios and non-IID data distributions [12, 25, 38, 46]. For example, FedCLIP [29] and PromptFL [11] directly extend the standard fine-tuning of CLIP to the federated setting to achieve strong performance. pFedprompt [10] combines a federated prompt learning scheme with personalized spatial visual features. Additionally, pFedPG [55] generates personalized prompts for each client based on their visual prompts to better align with their data distribution. FedOPT [21] utilizes knowledge from both personal and global textual prompts for prediction through unbalanced optimal transport. pFedMoAP [31] personalizes the prompt learning process through the lens of Mixture of Experts. However, previous methods have primarily focused on supervised federated learning with VLMs. In this paper, we explore leveraging VLMs for unsupervised federated classification, capitalizing on their zero-shot capabilities.

### 2.3. Unsupervised Learning for VLMs

In real-world applications, high annotation costs are often required to ensure that each data source has labeled data. This necessity drives us to develop effective methods for utilizing unlabeled data in downstream tasks. Pseudo labeling strategy [14, 16, 35, 47, 60] and entropy minimization [26] have been widely studied in this context. For instance, UPL [14] and FPL [35] select an equal number of pseudo labels for each category, while CPL [60] generates multiple candidate labels per instance to improve labeling accuracy. POUF [47] fine-tunes the model or

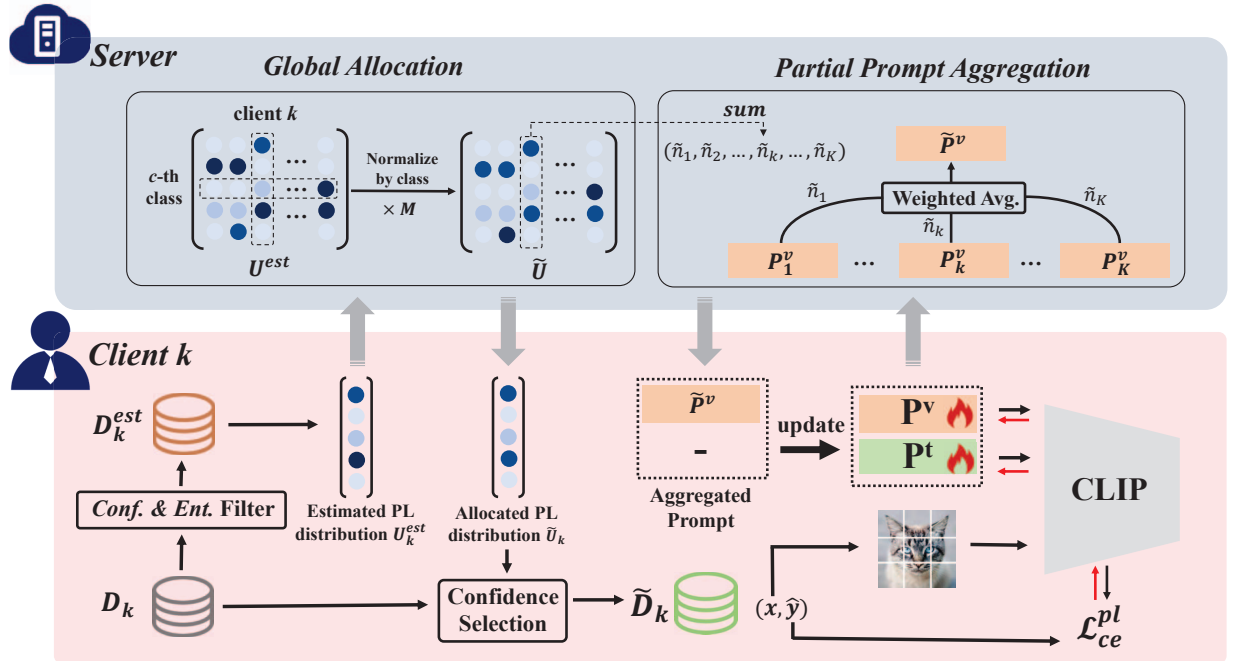


Figure 1. **The overview of our FedCoPL.** For pseudo labeling, FedCoPL begins by filtering the initial unlabeled samples to estimate the local distributions, which are uploaded to the server. Then, the server globally selects  $M$  pseudo labels for each category and allocates them to clients based on local estimated distributions. To handle label skews during training, we aggregate only visual prompts on the server to enhance global performance because the differences in textual prompts are significantly greater than those found in visual prompts.

prompt through the alignment of discrete label distributions between the source prompts and the unlabeled target data. UEO [26] adopts a confidence-aware strategy, minimizing the conditional entropy of high-confidence samples while simultaneously maximizing the marginal entropy of low-confidence ones. Despite their effectiveness, these pseudo labeling strategies face substantial challenges when directly applied to decentralized unsupervised federated settings [14]. To address this challenge, we propose a novel cooperative pseudo-labeling strategy that mitigates the inherent biases of pre-trained VLMs and significantly improves the quality of pseudo labels.

### 3. Method

In this section, we provide a detailed formulation of unsupervised federated learning for classification tasks with CLIP and introduce our proposed method. First, in Sec. 3.1, we review key principles underlying CLIP [39] and prompt tuning. Next, Sec. 3.2 outlines the framework of unsupervised federated learning with CLIP and presents its associated challenges. We then introduce our proposed method, FedCoPL, which is built upon two key strategies: the cooperative pseudo labeling strategy discussed in Sec. 3.3, and the partial prompt aggregation protocol covered in Sec. 3.4.

#### 3.1. Preliminary

We employ CLIP [39] as the foundational model for our framework. CLIP utilizes a dual-branch architecture consisting of an image encoder,  $f^v(\cdot)$ , and a text encoder,  $f^t(\cdot)$ , where each encoder processes data from its respective modality. For zero-shot predictions in downstream tasks, CLIP utilizes a human-designed prompt (e.g., “a photo of a [CLASS]”) for each class. Consider a  $C$ -way classification task as an example, textual embeddings of all classes  $\{f_c^t\}_{c=1}^C$  and the visual embedding of the test image  $f^v(x)$  are derived from the text and image encoders, respectively. The probability that image  $x$  belongs to the  $c$ -th category is calculated by applying the softmax operation as follows:

$$p_c(x) = \frac{\exp(\text{sim}(f^v(x), f_c^t)/\tau)}{\sum_{j=1}^C \exp(\text{sim}(f^v(x), f_j^t)/\tau)}, \quad (1)$$

where  $\tau$  is a temperature parameter. To enhance performance in downstream tasks, prompt tuning is widely adopted as a parameter-efficient fine-tuning method [63]. This involves introducing additional learnable textual tokens  $P^t$  and visual tokens  $P^v$  [15, 52] (referred to as textual and visual prompts) into the corresponding encoders, to adapt the original CLIP model to specific downstream tasks.

### 3.2. Unsupervised Federated Learning with CLIP

The powerful zero-shot classification capability of CLIP reduces reliance on labeled data, enabling the solution of previously intractable classification tasks in the context of UFL. Therefore, we extend the UFL paradigm to classification tasks by leveraging CLIP. In this setting, each client is initialized with the same pre-trained CLIP model and possesses an unlabeled local dataset  $D_k$  with a capacity of  $n_k$ . Additionally, each client is provided with the full set of class names for the classification task. Similar to standard federated learning frameworks, clients periodically upload non-data knowledge (e.g., statistical information, model parameters) to the server for aggregation and download global updates in return. To avoid ambiguity, we clarify that the label distribution and label skew referenced in this paper refer to the ground-truth labels of clients' datasets, which remain unknown during training.

The primary challenge arises from the inherent bias of CLIP and the data heterogeneity, particularly the label skew across clients. Specifically, while the union of all client datasets  $\bigcup_{k \in K} D_k$  encompasses samples from all categories, each individual dataset  $D_k$  may contain only a subset of these samples. Therefore, the direct application of CLIP within the UFL framework can result in an imbalanced training process across different categories. Additionally, the label skew among clients can lead to discrepancies in the model update directions across clients, which can hinder effective collaboration.

### 3.3. Cooperative Pseudo Labeling

To address the complex challenges inherent in unsupervised heterogeneous federated learning, we propose an effective and robust framework, FedCoPL. Our method incorporates two core strategies: a cooperative pseudo labeling (CoPL) strategy and a partial prompt aggregation protocol. A visual overview of the framework is provided in Fig. 1, with detailed pseudocode presented in the supplemental material.

Given the inherent bias of CLIP and the presence of label skews in the unlabeled data across clients, the direct application of conventional pseudo labeling methods can result in globally unbalanced and weak representative pseudo labels. Specifically, conventional pseudo labeling methods tend to select pseudo labels from categories favored by CLIP, even when these pseudo labels are incorrect [35, 60]. To address this, we introduce a cooperative pseudo labeling strategy to select a more representative pseudo label distribution and ensure a balanced distribution across all classes.

Specifically, for each client  $k$ , we filter samples from the original unlabeled set  $D_k$  according to model prediction confidence and entropy [14, 44] to construct the estimated set  $D_k^{est}$  as follows:

$$D_k^{est} = \{(x, \hat{y}) \mid \max_c p_c(x) > \tau_1, Ent(p(x)) < \tau_2\}, \quad (2)$$

where  $\hat{y} = \arg \max_c p_c(x)$  represents the predicted pseudo label, and  $Ent(\cdot)$  denotes the Shannon Entropy operator. Here,  $\tau_1$  and  $\tau_2$  represent the medians of the pseudo label confidences and the entropies of the training set, respectively. Due to the bias in CLIP,  $D_k^{est}$  tends to contain pseudo labels from categories favored by CLIP, even when these pseudo labels are incorrect [35, 60]. For example, suppose the true label distribution of  $D_k$  is [0.1, 0.1, 0.8], while the pseudo label distribution of  $D_k^{est}$  maybe [0.6, 0.05, 0.35], which has low accuracy and fails to represent the true label distribution. To address this issue, we propose a global allocation strategy that redistributes adjusted pseudo label distributions to clients in proportion to their locally estimated label distributions.

To make clients' pseudo label distributions more representative, we compute category-wise statistics on the filtered pseudo labels, yielding the estimated distribution  $U_k^{est}$ , a  $C$ -dimensional vector where the  $c$ -th element  $u_{k,c} = \sum_i \mathbb{1}(\hat{y}_i = c)$  indicates the number of samples associated with class  $c$  in client  $k$  according to the pseudo labels. Then, clients upload their estimated distribution  $U_k^{est}$  to the server for the collaborative allocation. Following [57, 60], we assume a uniform joint label distribution across all clients. To ensure adequate training for all classes, the server globally selects  $M$  pseudo labels for each category and allocates them to clients based on estimated distributions, as follows:

$$\begin{aligned} \tilde{U}_k &= (\tilde{u}_{k,1}, \tilde{u}_{k,2}, \dots, \tilde{u}_{k,C}), \\ \tilde{u}_{k,c} &= \left\lceil \frac{u_{k,c}}{\sum_{i \in K} u_{i,c}} \cdot M \right\rceil, \end{aligned} \quad (3)$$

where  $M = \frac{\sum_{k \in K} \sum_{c \in C} u_{k,c}}{C}$  is the global number of pseudo labels per class and  $\tilde{u}_{k,c}$  denotes the amount of pseudo labels of category  $c$  assigned to client  $k$ .

Finally, client  $k$  constructs its training set  $\tilde{D}_k$  based on the adjusted capacity  $\tilde{U}_k$ . Specifically, for each class  $c$ , the  $\tilde{u}_{k,c}$  samples with the highest prediction probabilities in the original dataset  $D_k$  are selected and incorporated into  $\tilde{D}_k$ . As described in Eq. (3), the overall sample budget  $M$  is allocated proportionally according to the clients' estimated distributions within each category, rather than relying on the absolute values of these estimates. This allocation strategy guarantees that, even when the estimated distributions are noisy or fail to fully capture the client's true local label distribution, the pseudo labels selected by our method remain both representative and accurate, thereby preserving the essential characteristics of the client's data and supporting robust collaborative learning.

As federated training progresses, we will periodically repeat these steps to generate, estimate, assign, and select pseudo labels. This iterative procedure ensures that the

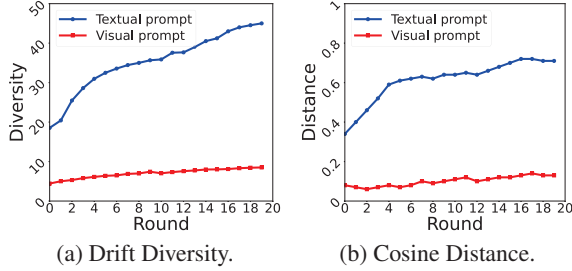


Figure 2. Drift diversity and cosine distance of prompts among clients during training in CIFAR10 [19] dataset. The differences observed in textual prompts are significantly greater than those observed in visual prompts.

pseudo-labeled training data  $\tilde{D}_k$  is regularly updated, maintaining a globally balanced and high-accuracy dataset.

### 3.4. Partial Prompt Aggregation

For each client, we select both textual prompt  $P_k^t$  and visual prompt  $P_k^v$  to be optimized using the local training set  $\tilde{D}_k$ . Since each client’s training set pertains to the same task, effective knowledge aggregation can enhance the overall framework’s performance. A straightforward approach is to aggregate the prompts using a simple averaging operation. However, due to label skews among clients, this may lead to suboptimal or even detrimental performance.

As for CLIP, updates to the visual branch primarily enhance image representation knowledge, while the textual branch focuses on determining the classification boundaries by leveraging category information [52]. Consequently, under heterogeneous data distributions among clients, where clients share the same task but different label distributions, the visual prompts  $\{P_k^v\}_{k=1}^K$  are more likely to be similar, while the textual prompts  $\{P_k^t\}_{k=1}^K$  exhibit greater variability. To validate this conjecture, we measure the differences in both visual and textual prompts across all clients using drift diversity [20] and cosine distance, which capture the magnitude and directional differences in prompt updates, respectively. As illustrated in Fig. 2, the differences in textual prompts are significantly greater than those in visual prompts, which confirms our hypothesis.

To avoid conflicts when aggregating widely different textual prompts and to promote cooperation among clients, we propose a partial prompt aggregation protocol. In this protocol, all clients upload their visual prompts  $\{P_k^v\}_{k=1}^K$  to the server for aggregation while keeping their textual prompts  $\{P_k^t\}_{k=1}^K$  locally for personalization. For the server, we design the visual prompt aggregation strategy that utilizes a weighted average approach following [34] as follows:

$$P^v = \sum_{k=1}^K \frac{\tilde{n}_k}{\sum_{i \in K} \tilde{n}_i} \cdot P_k^v, \quad (4)$$

where  $\tilde{n}_k = \sum_c \tilde{u}_{k,c}$  represents the total number of samples assigned to client  $k$ . The aggregated visual prompts  $P^v$  are then distributed back to the clients as initialization for their local models for the next training round. The overall objective function of FedCoPL is formulated as follows:

$$\min_{P^v, \{P_k^t\}_{k=1}^K} \sum_{k=1}^K \mathbb{E}_{(x, \hat{y}) \in \tilde{D}_k} \ell_{ce}(g(P^v, P_k^t; x), \hat{y}), \quad (5)$$

where  $g(P^v, P_k^t; \cdot)$  is the model output,  $\hat{y}$  denotes pseudo labels, and  $\ell_{ce}(\cdot, \cdot)$  is the cross-entropy loss function.

This approach enhances global performance by aggregating visual prompts while allowing clients to utilize personalized textual prompts that align better with their specific data distributions. Moreover, in comparison to methods that aggregate prompts from both modalities, our approach can reduce communication overhead by aggregating only textual prompts, making it more practical for environments where communication resources are constrained.

**Remark.** The proposed FedCoPL framework preserves client privacy by ensuring that the uploaded estimated distributions are derived solely from CLIP’s predictions, making it infeasible to inversely infer the clients’ exact label distributions. This highlights the enhanced privacy-preserving nature of our method, especially when compared with several existing approaches [18, 61] that require the transmission of exact category distributions, potentially exposing sensitive client-specific information.

## 4. Experiments

### 4.1. Setups

**Datasets.** We evaluated the performance of our method on six public benchmark datasets characterized by varying types of label skew. Following previous research [5, 10, 21], we employed four representative visual classification datasets: DTD [4], RESISC45 [3], UCF101 [45], and CUB [49], along with two standard federated classification benchmark datasets: CIFAR10 and CIFAR100 [19]. We partitioned each dataset into distinct training and test sets and further split them into non-overlapping subsets for different clients. Specifically, we followed the settings outlined in [23] and employed two prevalent forms of label skew: quantity-based and Dirichlet-based label skews. In the quantity-based label skew, all training data is first grouped by label and then partitioned into shards of varying sizes. In the Dirichlet-based label skew, clients receive samples for each class based on the Dirichlet distribution [64]. Here, the parameter  $\beta$  controls the degree of label skew, with lower values indicating severe label skew.

**Baseline methods.** In our experiments, we compare FedCoPL, with two popular pseudo label selection methods in central unsupervised learning and four supervised fed-

Table 1. Accuracy (%) of different methods under Dirichlet-based label skews with ViT/B-32 as the backbone. FPL [35] and CPL [60] are adopted as baseline pseudo labeling (PL) method. **Bold** and underline represent the best and second-best results.

Method	PL	DTD		RESISC45		CUB		UCF101		CIFAR10		CIFAR100		Average	
		$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.1$	$\beta = 0.05$
CLIP	-	43.24	43.24	54.51	54.51	<u>51.28</u>	<u>51.28</u>	61.00	61.00	86.93	86.93	64.17	64.17	60.19	60.19
PromptFL [11]	FPL	45.79	44.62	59.76	58.05	47.29	46.16	64.39	62.96	87.51	85.02	63.26	62.56	61.33	59.90
PromptFL [11]	CPL	44.84	46.32	62.52	60.99	48.72	48.86	63.86	64.97	88.90	89.14	66.04	65.14	62.48	62.57
Promptprox [24]	FPL	45.15	43.88	59.36	58.59	47.04	47.01	62.91	61.27	87.13	85.62	62.77	64.12	60.73	60.08
Promptprox [24]	CPL	43.51	45.85	59.99	62.35	49.25	48.65	64.55	63.94	90.44	90.07	65.86	66.15	62.27	62.84
pFedPrompt [10]	FPL	44.56	46.19	<u>65.95</u>	60.52	48.48	44.42	64.37	64.18	87.45	90.72	65.08	65.83	62.65	61.98
pFedPrompt [10]	CPL	44.22	47.59	61.76	<u>66.79</u>	47.23	50.97	<u>65.59</u>	<u>65.44</u>	90.26	89.10	65.63	<u>67.86</u>	62.45	<u>64.63</u>
FedOPT [21]	FPL	<u>51.97</u>	<u>48.71</u>	57.76	55.62	47.99	47.05	61.83	64.56	<u>91.32</u>	90.73	<u>67.51</u>	66.50	<u>63.06</u>	62.20
FedOPT [21]	CPL	37.43	39.33	51.10	48.47	46.61	45.07	58.93	57.12	91.08	<u>91.36</u>	57.87	59.81	57.17	56.86
<b>FedCoPL</b>	<b>CoPL</b>	<b>60.89</b>	<b>66.37</b>	<b>75.76</b>	<b>80.26</b>	<b>56.09</b>	<b>54.80</b>	<b>73.20</b>	<b>74.97</b>	<b>95.38</b>	<b>95.11</b>	<b>73.59</b>	<b>72.84</b>	<b>72.49</b>	<b>74.06</b>

Table 2. Accuracy (%) of different methods under quantity-based label skews with ViT/B-32 as the backbone. FPL [35] and CPL [60] are adopted as baseline pseudo labeling (PL) method. **Bold** and underline represent the best and second-best results.

Method	PL	DTD	RESISC45	CUB	UCF101	CIFAR10	CIFAR100	Average
CLIP	-	43.24	54.51	<u>51.28</u>	61.00	86.93	64.17	60.19
PromptFL [11]	FPL	43.53	57.53	46.66	62.35	88.29	63.94	60.38
PromptFL [11]	CPL	43.82	61.11	47.31	64.23	88.70	<u>67.23</u>	62.07
Promptprox [24]	FPL	44.89	56.71	48.74	61.80	85.69	64.01	60.31
Promptprox [24]	CPL	45.53	60.77	47.61	63.86	89.21	66.49	62.25
pFedPrompt [10]	FPL	45.09	60.46	47.71	65.39	87.54	62.82	61.50
pFedPrompt [10]	CPL	45.14	<u>66.68</u>	50.45	<u>65.57</u>	87.98	65.45	<u>63.55</u>
FedOPT [21]	FPL	<u>47.65</u>	56.04	46.61	62.95	<u>90.74</u>	66.76	61.79
FedOPT [21]	CPL	40.24	47.09	44.00	57.24	89.21	57.98	55.96
<b>FedCoPL</b>	<b>CoPL</b>	<b>56.18</b>	<b>81.06</b>	<b>56.31</b>	<b>72.03</b>	<b>94.86</b>	<b>73.39</b>	<b>72.31</b>

erated learning methods. Regarding pseudo labeling strategy, **FPL** [35] selects the most reliable samples based on confidence for each class, while **CPL** [60] generates multiple pseudo labels for each sample through selection at both the sample and category levels. Zero-shot prediction involves utilizing the pre-trained CLIP model with a hand-crafted textual prompt template, such as “a photo of a [CLASS],” to predict the test data. For federated training methods, **PromptFL** [11] employs shared prompts learned across clients via FedAvg [34]. **Promptprox**, introduced in [11], is derived from the traditional federated learning technique FedProx [24], which employs the proximal term to constrain the update direction of the local model. **pFed-Prompt** [10] learns a unified textual prompt with personalized attention modules for local visual embeddings. Finally, **FedOPT** [21] performs conditional optimal transport between the global and local textual prompts to effectively fuse global and local knowledge.

**Implementation details.** All results in this paper are based on a frozen CLIP using two backbones, ViT-B/16 [6] and ViT-B/32 [6], with ViT-B/32 as the default unless otherwise specified. We set the number of clients,  $K$ , to 5 for the CUB [49] and UCF101 [45] datasets, and to 10 for the other datasets, with full client participation as the default setting.

We conduct 20 communication rounds for all experimental datasets, where pseudo labels are updated every 5 rounds using the latest local model. Within each communication round, each client performs 10 epochs of local training. We optimize the prompts using mini-batch Stochastic Gradient Descent (SGD) with a learning rate of 0.1, a momentum of 0.9, and decay following the cosine annealing rule. We conduct three trials for each experimental setting and report the mean accuracy. All experiments are conducted using PyTorch [36] on NVIDIA 3090 GPUs. More details on the implementation can be found in the supplementary material.

## 4.2. Experimental Results

**Results under Dirichlet-based label skews with various datasets.** Table 1 reports the performance of different methods under Dirichlet-based label skews ( $\beta \in 0.1, 0.05$ ). Our method significantly outperforms state-of-the-art algorithms across all datasets, demonstrating the effectiveness of our cooperative pseudo labeling strategy and partial prompt aggregation protocol. Notably, FedOPT [21], the latest personalized supervised federated prompt tuning method, shows that integrating it with baseline pseudo label selection methods leads to worse performance than zero-shot learning. This is attributed to the low accuracy of the pseudo labels chosen by each client, which fail to capture

Table 3. Accuracies (%) of combining our pseudo labeling strategy CoPL with existing federated training methods under Dirichlet-based label skews ( $\beta = 0.1$ ). FPL [35] is adopted as baseline pseudo labeling (PL) method. (The values) represent the performance gains.

Method	PL	DTD	RESISC45	CUB	UCF101	CIFAR10
PromptFL [11]	FPL	45.79	59.76	47.29	64.39	87.51
PromptFL [11]	<b>CoPL</b>	47.34 (+1.55)	61.68 (+1.92)	49.08 (+1.79)	64.97 (+0.58)	90.20 (+2.69)
pFedPrompt [10]	FPL	44.56	65.95	48.48	64.37	87.45
pFedPrompt [10]	<b>CoPL</b>	49.25 (+4.69)	67.97 (+2.02)	52.77 (+2.29)	67.93 (+3.56)	91.13 (+3.68)
FedOPT [21]	FPL	51.97	57.76	47.99	61.83	91.32
FedOPT [21]	<b>CoPL</b>	57.93 (+5.96)	74.85 (+17.09)	55.72 (+7.73)	69.21 (+7.38)	94.02 (+2.70)

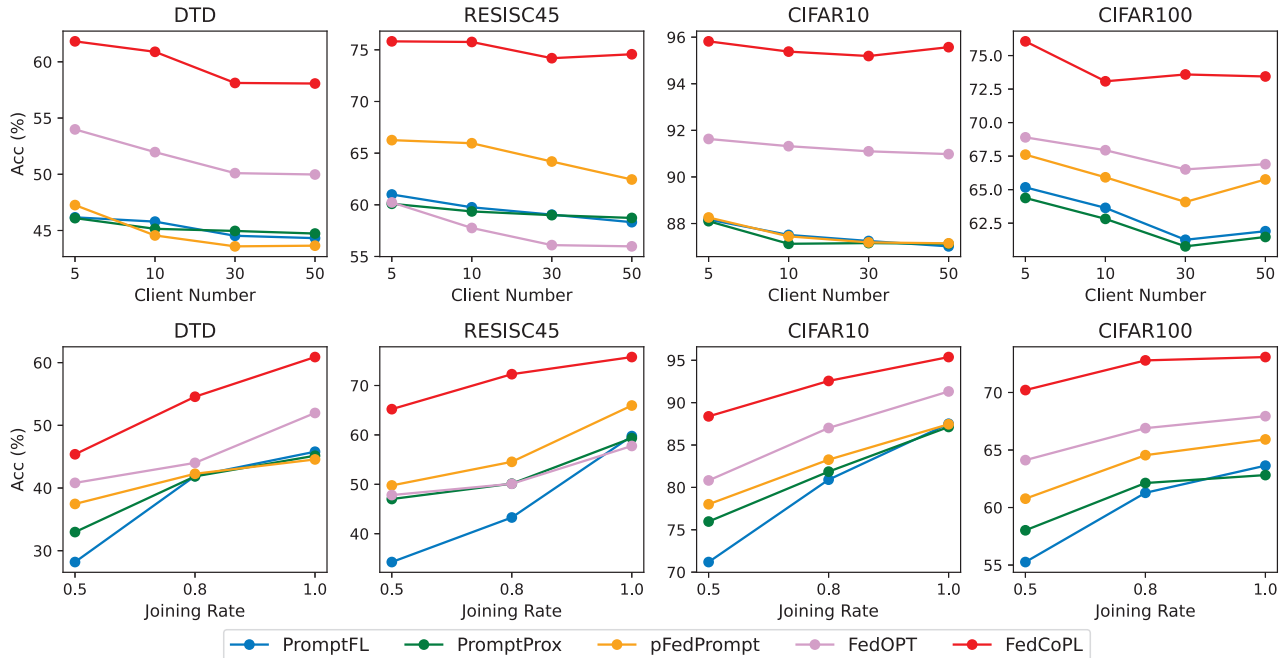


Figure 3. Results of experiments with various client numbers and different client joining rates under Dirichlet-based label skews ( $\beta = 0.1$ ). FPL [35] is adopted as the baseline pseudo labeling method.

the true local data distribution, and this issue is further exacerbated by personalized training. A more detailed discussion is provided in the following sections. In the CUB [49] dataset, all baseline methods perform worse than zero-shot prediction owing to the worse representativeness of selected pseudo labels. In contrast, our method consistently achieves superior performance on all datasets, underscoring the importance of the cooperative pseudo labeling strategy and the partial prompt aggregation protocol.

**Results under quantity-based label skews with various datasets.** We present the performance of all methods under quantity-based label distribution skews in Table 2, where the parameter is set as  $s = C \times 0.2$  across all datasets, with  $C$  representing the number of classes in each dataset. In this setting, each client possesses samples from only a few classes, which poses challenges for pseudo labeling and model training. Owing to this, many baseline results are outperformed by the zero-shot approach. For example, Fe-

dOPT with CPL performs 7.42% and 7.28% worse than the zero-shot accuracy on the RESISC45 and CUB datasets, respectively. In contrast, our method maintains strong performance, similar to that observed with Dirichlet-based distributions, further highlighting the effectiveness of FedCoPL.

### 4.3. Analysis

**Comparison of different pseudo labeling methods.** We compare our pseudo labeling method with two baseline methods, FPL [35] and CPL [60]. As shown in Table 1 and Table 2, for the same federated training method, CPL outperforms FPL on datasets such as DTD [4] and CUB [49], whereas the opposite is true for the CIFAR datasets. In contrast, our proposed cooperative pseudo labeling strategy consistently outperforms both baseline methods across all datasets. To further demonstrate the effectiveness of our approach, we integrate our pseudo labeling method into various state-of-the-art supervised federated training frame-

Table 4. Experimental results (%) using CLIP ViT-B/16 as backbone under Dirichlet-based label skew ( $\beta = 0.1$ ). FPL [35] is adopted as the baseline pseudo labeling (PL) method.

Method	PL	DTD	RESISC45	CUB	UCF101
Zero-shot CLIP	-	42.87	56.61	55.16	65.13
PromptFL [11]	FPL	44.36	61.36	51.95	64.60
PromptProx [24]	FPL	47.12	61.01	49.44	65.31
pFedPrompt [10]	FPL	47.65	63.18	50.19	65.19
FedOPT [21]	FPL	46.52	58.94	49.67	65.68
<b>FedCoPL</b>	<b>CoPL</b>	<b>56.83</b>	<b>77.50</b>	<b>59.78</b>	<b>78.29</b>

works. As shown in Table 3, our method significantly enhances the performance of baseline pseudo label selection methods across different datasets, with a notable increase of up to 17.09% on the RESISC45 dataset. The results clearly validate the effectiveness of our pseudo labeling method compared to baseline methods.

**Results under different client numbers.** We analyze the performance of our proposed method compared to baseline methods under varying clients numbers. Unless stated otherwise, our experiments are conducted under Dirichlet-based skews with  $\beta = 0.1$ . We partition the DTD, RESISC45 and CIFAR datasets among 5, 10, 30, and 50 clients and report their final accuracy in the first row of Figure 3. Notably, as the number of clients increases, the performance of baseline methods declines significantly, whereas our method consistently maintains high accuracy. This underscores the robustness of our approach, demonstrating its capability to maintain effective learning performance even under varying numbers of participating clients.

**Impact of client joining rates.** In this analysis, we investigate the impact of variations in client participation rates, considering values from 0.5, 0.8, 1.0. As illustrated in the second row of Figure 3, our approach consistently outperforms competing methods across all participation rates. As the client participation rate decreases, the performance of all methods deteriorates significantly. This instability is anticipated, as a lower participation rate exacerbates the divergence between randomly selected clients and the global model, resulting in unstable convergence. Nevertheless, our approach maintains superior performance, demonstrating its robustness to fluctuations in client participation rates.

**Results under different image encoder backbone.** We further conduct experiments to evaluate the effect of different image encoders. The comparison results using ViT-B/16 are presented in Table 4. Our approach consistently outperforms previous methods, highlighting the effectiveness of our strategy in improving the performance of smaller image encoders. Additional results using RN50 as the image encoder backbone can be found in the supplementary material. These experiments emphasize the robustness of FedCoPL in real-world FL scenarios, particularly when clients have limited computational resources.

Table 5. Ablation study. Accuracies (%) under Dirichlet-based label skews. Conf. and Ent. denote confidence-based and entropy-based filters. G.A. is the global allocation of pseudo labels.

Conf.	Ent.	G.A.	DTD	RESISC45	CUB	UCF101
-	-	-	45.79	59.76	47.29	64.39
-	-	✓	46.35	69.60	51.98	65.31
-	✓	✓	58.83	72.28	54.91	70.23
✓	-	✓	55.59	73.66	50.38	69.13
✓	✓	✓	60.89	75.76	56.09	73.20

Table 6. Ablation results (%) comparing the effects of aggregating textual prompts (T.P.) only, visual prompts (V.P.) only, both, or none under Dirichlet-based label skew ( $\beta = 0.1$ ).

T.P.	V.P.	DTD	RESISC45	CUB	UCF101
-	-	50.45	61.72	49.08	66.02
✓	✓	47.34	61.68	49.08	64.97
✓	-	55.31	67.12	51.76	69.53
-	✓	60.89	75.76	56.09	73.20

**Effectiveness of each component.** Our approach consists of two key modules: a cooperative pseudo labeling strategy and a partial prompt aggregation protocol. The results presented in Table 5 demonstrate the effectiveness of the two filtering criteria and global allocation. To further validate the effectiveness of our partial prompt aggregation, we conduct an additional experiment, as shown in Table 6. Compared to aggregating two modal prompts, aggregating textual prompts only, or retaining two modal prompts locally, the partial prompt aggregation yields the optimal results, demonstrating its effectiveness.

## 5. Conclusion

In this paper, we extend unsupervised federated learning to classification tasks using CLIP for the first time. In this setting, clients with pre-trained CLIP models and unlabeled data collaborate in training to enhance performance without data sharing. To address the challenge of CLIP’s internet bias and potential label skew among clients, we propose FedCoPL, which includes a cooperative pseudo labeling strategy and a partial prompt aggregation protocol. Specifically, each client estimates and uploads its pseudo label distribution, and the server adjusts and redistributes them to mitigate global imbalance across categories. Moreover, the aggregation protocol aggregates only visual prompts on the server to improve global performance, while textual prompts are kept locally for better personalization by each client. Extensive experimental results demonstrate that our FedCoPL outperforms baseline methods across various degrees of label skew. In future work, we will conduct a theoretical analysis of FedCoPL, including convergence, privacy, fairness, and other pertinent considerations.

## Acknowledgments

This work was funded by the National Natural Science Foundation of China under Grants (62276256, U2441251), the Young Elite Scientists Sponsorship Program by CAST (2023QNRC001), and the Young Scientists Fund of the State Key Laboratory of Multimodal Artificial Intelligence Systems (ES2P100117).

## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Proc. NeurIPS*, 2022. 1
- [2] Debora Caldarola, Pietro Cagnasso, Barbara Caputo, and Marco Ciccone. Beyond local sharpness: Communication-efficient global sharpness-aware minimization for federated learning. In *Proc. CVPR*, 2025. 1
- [3] Gong Cheng, Junwei Han, and Xiaoqiang Lu. Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE*, 2017. 5
- [4] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proc. CVPR*, 2014. 5, 7
- [5] Tianyu Cui, Hongxia Li, Jingya Wang, and Ye Shi. Harmonizing generalization and personalization in federated prompt learning. In *Proc. ICML*, 2024. 5
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [7] Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Céline Hudelot, and Pierre Colombo. Colpali: Efficient document retrieval with vision language models. In *Proc. ICLR*, 2024. 1
- [8] Kuangpu Guo, Yuhe Ding, Jian Liang, Ran He, Zilei Wang, and Tieniu Tan. Exploring vacant classes in label-skewed federated learning. In *Proc. AAAI*, 2025. 1
- [9] Shunxin Guo, Hongsong Wang, and Xin Geng. Dynamic heterogeneous federated learning with multi-level prototypes. *Pattern Recognition*, 153:110542, 2024. 1
- [10] Tao Guo, Song Guo, and Junxiao Wang. Pfdprompt: Learning personalized prompt for vision-language models in federated learning. In *Proceedings of the ACM Web Conference 2023*, 2023. 2, 5, 6, 7, 8
- [11] Tao Guo, Song Guo, Junxiao Wang, Xueyang Tang, and Wenchao Xu. Promptfl: Let federated participants cooperatively learn prompts instead of models—federated learning in age of foundation model. *IEEE Transactions on Mobile Computing*, 23(5):5179–5194, 2023. 2, 6, 7, 8
- [12] Shaunak Halbe, James Seale Smith, Junjiao Tian, and Zsolt Kira. Hepco: Data-free heterogeneous prompt consolidation for continual federated learning. *arXiv preprint arXiv:2306.09970*, 2023. 2
- [13] Sungwon Han, Sungwon Park, Fangzhao Wu, Sundong Kim, Chuhan Wu, Xing Xie, and Meeyoung Cha. Fedx: Unsupervised federated learning with cross knowledge distillation. In *Proc. ECCV*, 2022. 1
- [14] Tony Huang, Jack Chu, and Fangyun Wei. Unsupervised prompt learning for vision-language models. *arXiv preprint arXiv:2204.03649*, 2022. 1, 2, 3, 4
- [15] Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual prompt tuning. In *Proc. ECCV*, 2022. 3
- [16] Yunpeng Jia, Xiufen Ye, Xinkui Mei, Yusong Liu, and Shuxiang Guo. Mltu: mixup long-tail unsupervised zero-shot image classification on vision-language models. *Multimedia Systems*, 30(3):169, 2024. 2
- [17] Hansol Kim, Youngjun Kwak, Minyoung Jung, Jinho Shin, Youngsung Kim, and Changick Kim. Protofl: Unsupervised federated learning via prototypical distillation. In *Proc. ICCV*, 2023. 1
- [18] Woojung Kim, Keondo Park, Kihyuk Sohn, Raphael Shu, and Hyung-Sin Kim. Federated semi-supervised learning with prototypical networks. *arXiv preprint arXiv:2205.13921*, 2022. 5
- [19] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [20] Bo Li, Mikkel N Schmidt, Tommy S Alstrøm, and Sebastian U Stich. On the effectiveness of partial variance reduction in federated learning with heterogeneous data. In *Proc. CVPR*, 2023. 5
- [21] Hongxia Li, Wei Huang, Jingya Wang, and Ye Shi. Global and local prompts cooperation via optimal transport for federated learning. In *Proc. CVPR*, 2024. 2, 5, 6, 7, 8
- [22] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 1
- [23] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *Proc. ICDE*, 2022. 1, 5
- [24] 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, 8
- [25] Xi Li and Jiaqi Wang. Position paper: Assessing robustness, privacy, and fairness in federated learning integrated with foundation models. *arXiv preprint arXiv:2402.01857*, 2024. 2
- [26] Jian Liang, Lijun Sheng, Zhengbo Wang, Ran He, and Tieniu Tan. Realistic unsupervised clip fine-tuning with universal entropy optimization. In *Proc. ICML*, 2024. 2, 3
- [27] Xinting Liao, Weiming Liu, Chaochao Chen, Pengyang Zhou, Fengyuan Yu, Huabin Zhu, Binhui Yao, Tao Wang, Xiaolin Zheng, and Yanchao Tan. Rethinking the representation in federated unsupervised learning with non-iid data. In *Proc. CVPR*, 2024. 1, 2
- [28] Nan Lu, Zhao Wang, Xiaoxiao Li, Gang Niu, Qi Dou, and Masashi Sugiyama. Federated learning from only unlabeled data with class-conditional-sharing clients. In *Proc. ICLR*, 2022. 2

- [29] Wang Lu, Xixu Hu, Jindong Wang, and Xing Xie. Fedclip: Fast generalization and personalization for clip in federated learning. *arXiv preprint arXiv:2302.13485*, 2023. 2
- [30] Ekdeep Lubana, Chi Ian Tang, Fahim Kawsar, Robert Dick, and Akhil Mathur. Orchestra: Unsupervised federated learning via globally consistent clustering. In *Proc. ICML*, 2022. 2
- [31] Jun Luo, Chen Chen, and Shandong Wu. Mixture of experts made personalized: Federated prompt learning for vision-language models. *arXiv preprint arXiv:2410.10114*, 2024. 2
- [32] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. In *Proc. NeurIPS*, 2021. 1
- [33] Ana-Maria Marcu, Long Chen, Jan Hünemann, Alice Karnsund, Benoit Hanotte, Prajwal Chidananda, Saurabh Nair, Vijay Badrinarayanan, Alex Kendall, Jamie Shotton, et al. Lingoqa: Visual question answering for autonomous driving. In *Proc. ECCV*, 2024. 1
- [34] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proc. AISTATS*, 2017. 5, 6
- [35] Cristina Menghini, Andrew Delworth, and Stephen Bach. Enhancing clip with clip: Exploring pseudolabeling for limited-label prompt tuning. In *Proc. NeurIPS*, 2023. 1, 2, 4, 6, 7, 8
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Proc. NeurIPS*, 2019. 6
- [37] Farhad Pourpanah, Mahdiyar Molahasani, Milad Soltany, Michael Greenspan, and Ali Etemad. Federated unsupervised domain generalization using global and local alignment of gradients. In *Proc. AAAI*, 2025. 1, 2
- [38] Chen Qiu, Xingyu Li, Chaithanya Kumar Mummadi, Madan Ravi Ganesh, Zhenzhen Li, Lu Peng, and Wan-Yi Lin. Text-driven prompt generation for vision-language models in federated learning. *arXiv preprint arXiv:2310.06123*, 2023. 2
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 2021. 1, 3
- [40] Yasar Abbas Ur Rehman, Yan Gao, Pedro Porto Buarque De Gusmão, Mina Alibeigi, Jiajun Shen, and Nicholas D Lane. L-dawa: Layer-wise divergence aware weight aggregation in federated self-supervised visual representation learning. In *Proc. ICCV*, 2023. 1, 2
- [41] Micah J Sheller, Brandon Edwards, G Anthony Reina, Jason Martin, Sarthak Pati, Aikaterini Kotrotsou, Mikhail Milchenko, Weilin Xu, Daniel Marcus, Rivka R Colen, et al. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*, 10(1):12598, 2020. 1
- [42] Yujun Shi, Jian Liang, Wenqing Zhang, Vincent Tan, and Song Bai. Towards understanding and mitigating dimensional collapse in heterogeneous federated learning. In *Proc. ICLR*, 2023. 1
- [43] Yujun Shi, Jian Liang, Wenqing Zhang, Chuhui Xue, Vincent YF Tan, and Song Bai. Understanding and mitigating dimensional collapse in federated learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5): 2936–2949, 2023. 1
- [44] Manli Shu, Weili Nie, De-An Huang, Zhiding Yu, Tom Goldstein, Anima Anandkumar, and Chaowei Xiao. Test-time prompt tuning for zero-shot generalization in vision-language models. In *Proc. NeurIPS*, 2022. 4
- [45] K Soomro. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5, 6
- [46] Shangchao Su, Mingzhao Yang, Bin Li, and Xiangyang Xue. Federated adaptive prompt tuning for multi-domain collaborative learning. In *Proc. AAAI*, 2024. 2
- [47] Korawat Tanwisuth, Shujian Zhang, Huangjie Zheng, Pengcheng He, and Mingyuan Zhou. Pouf: Prompt-oriented unsupervised fine-tuning for large pre-trained models. In *Proc. ICML*, 2023. 2
- [48] Ye Tian, Haolei Weng, and Yang Feng. Towards the theory of unsupervised federated learning: Non-asymptotic analysis of federated em algorithms. In *Proc. ICML*, 2024. 1
- [49] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. 5, 6, 7
- [50] Song Wen, Guian Fang, Renrui Zhang, Peng Gao, Hao Dong, and Dimitris Metaxas. Improving compositional text-to-image generation with large vision-language models. *arXiv preprint arXiv:2310.06311*, 2023. 1
- [51] Chulin Xie, De-An Huang, Wenda Chu, Daguang Xu, Chaowei Xiao, Bo Li, and Anima Anandkumar. Perada: Parameter-efficient federated learning personalization with generalization guarantees. In *Proc. CVPR*, 2024. 1
- [52] Yinghui Xing, Qirui Wu, De Cheng, Shizhou Zhang, Guoqiang Liang, Peng Wang, and Yanning Zhang. Dual modality prompt tuning for vision-language pre-trained model. In *Proc. ACM MM*, 2023. 2, 3, 5
- [53] Liao Xinting, Liu Weiming, Qian Jiaming, Zhou Pengyang, Xu Jiahe, Wang Wenjie, Chen Chaochao, Zheng Xiaolin, and Tat-Seng Chua. Focoop: Enhancing out-of-distribution robustness in federated prompt learning for vision-language models. In *Proc. CVPR*, 2025. 1
- [54] Hao Yan and Yuhong Guo. Lightweight unsupervised federated learning with pretrained vision language model. In *International Workshop on Trustworthy Federated Learning*, 2024. 1
- [55] Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient model personalization in federated learning via client-specific prompt generation. In *Proc. ICCV*, 2023. 2
- [56] Xiyuan Yang, Wenke Huang, and Mang Ye. Fedas: Bridging inconsistency in personalized federated learning. In *Proc. CVPR*, 2024. 1

- [57] Rui Ye, Mingkai Xu, Jianyu Wang, Chenxin Xu, Siheng Chen, and Yanfeng Wang. Feddisco: Federated learning with discrepancy-aware collaboration. In *Proc. ICML*, 2023. 4
- [58] Fengda Zhang, Kun Kuang, Long Chen, Zhaoyang You, Tao Shen, Jun Xiao, Yin Zhang, Chao Wu, Fei Wu, Yueting Zhuang, et al. Federated unsupervised representation learning. *Frontiers of Information Technology & Electronic Engineering*, 24(8):1181–1193, 2023. 1
- [59] Jianqing Zhang, Yang Hua, Hao Wang, Tao Song, Zhengui Xue, Ruhui Ma, and Haibing Guan. Fedala: Adaptive local aggregation for personalized federated learning. In *Proc. AAAI*, 2023. 1
- [60] Jiahua Zhang, Qi Wei, Feng Liu, and Lei Feng. Candidate pseudolabel learning: Enhancing vision-language models by prompt tuning with unlabeled data. In *Proc. ICML*, 2024. 2, 4, 6, 7
- [61] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proc. CVPR*, 2022. 5
- [62] Yuguang Zhang, Kuangpu Guo, Zhihe Lu, Yunbo Wang, and Jian Liang. Personalized federated learning via dual-prompt optimization and cross fusion. *arXiv preprint arXiv:2506.21144*, 2025. 1
- [63] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 3
- [64] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021. 5
- [65] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. Collaborative unsupervised visual representation learning from decentralized data. In *Proc. ICCV*, 2021. 2
- [66] Weiming Zhuang, Yonggang Wen, and Shuai Zhang. Divergence-aware federated self-supervised learning. In *Proc. CVPR*, 2022. 2