

# Exploring Vacant Classes in Label-Skewed Federated Learning

Kuangpu Guo<sup>1,2</sup>, Yuhe Ding<sup>3</sup>, Jian Liang<sup>2,4\*</sup>, Zilei Wang<sup>1</sup>, Ran He<sup>2,4</sup>, Tieniu Tan<sup>5,2,4</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>Anhui University

<sup>4</sup>University of Chinese Academy of Sciences

<sup>5</sup>Nanjing University

gkp@mail.ustc.edu.cn, {madao3c, liangjian92}@gmail.com, zlwang@ustc.edu.cn, {rhe, tnt}@nlpr.ia.ac.cn

## Abstract

Label skews, characterized by disparities in local label distribution across clients, pose a significant challenge in federated learning. As minority classes suffer from worse accuracy due to overfitting on local imbalanced data, prior methods often incorporate class-balanced learning techniques during local training. Although these methods improve the mean accuracy across all classes, we observe that vacant classes—referring to categories absent from a client’s data distribution—remain poorly recognized. Besides, there is still a gap in the accuracy of local models on minority classes compared to the global model. This paper introduces FedVLS, a novel approach to label-skewed federated learning that integrates both vacant-class distillation and logit suppression simultaneously. Specifically, vacant-class distillation leverages knowledge distillation during local training on each client to retain essential information related to vacant classes from the global model. Moreover, logit suppression directly penalizes network logits for non-label classes, effectively addressing misclassifications in minority classes that may be biased toward majority classes. Extensive experiments validate the efficacy of FedVLS, demonstrating superior performance compared to previous state-of-the-art (SOTA) methods across diverse datasets with varying degrees of label skews.

**Code** — <https://github.com/krumpguo/FedVLS>

## Introduction

Federated learning has emerged as a prominent distributed learning paradigm, lauded for its capability to train a global model without direct access to raw data (Konečný et al. 2016; Li et al. 2020; Kairouz et al. 2021). The traditional federated learning (FL) algorithm, FedAvg (McMahan et al. 2017), follows an iterative process of refining the global model by aggregating parameters from local models, which are initialized with the latest global model parameters and trained across diverse client devices (Sheller et al. 2020; Li et al. 2020; Chai et al. 2024; Luo et al. 2022). In real-world scenarios, local client data often originate from diverse populations or organizations, displaying significant label skews

that severely undermine the performance of federated learning (Hsu, Qi, and Brown 2019; Yang, Fang, and Liu 2021; Reguieg et al. 2023; Zhang et al. 2023; Ye et al. 2023).

Typically, the local data often consists of majority classes and minority classes, which refer to classes with a large amount of data a small amount of data, respectively (Zhang et al. 2022, 2023). As evidenced in prior studies (Zhang et al. 2022; Chen et al. 2022), the accuracy of minority classes notably decreases after local updates, signaling that the client model is overfitting to local imbalanced data. Consequently, this results in substantial performance degradation of the global model (Yeganeh et al. 2020; Li et al. 2020; Liu et al. 2022). To address the issue of lower accuracy in minority classes, previous methods often incorporate class-balanced learning techniques during local training (Zhang et al. 2022; Chen et al. 2022; Wang et al. 2023b; Shen, Wang, and Lv 2023). Some works (Zhang et al. 2022; Chen et al. 2022; Shen, Wang, and Lv 2023; Wang et al. 2023b) advocate for calibrating logits according to the client data distribution to balance minority and majority classes. However, previous methods have overlooked vacant classes, which refer to classes without data but have highly versatile applications. For instance, in landmark detection (Weyand et al. 2020), most contributors possess only a subset of landmark categories from places they have lived or traveled. More importantly, these vacant classes can significantly compromise the model’s performance, particularly in scenarios with highly skewed label distributions.

For example, we compare the class-wise accuracy of the initial global model and updated local models using both the classic method FedAvg (McMahan et al. 2017) and one SOTA method FedLC (Zhang et al. 2022). As depicted in Figure 1 (b) and (c), the updated local model exhibits a notable decline in accuracy for vacant classes (e.g., categories 0, 1, 2, 4, 6, and 7) compared to the class-wise accuracy of the initial global model. In extreme cases, the accuracy even decreases close to zero (e.g., category 6). We posit this severe decline is attributed to the loss of information about vacant classes in the updated local models. By the way, although FedLC (Zhang et al. 2022) partially alleviates the performance decline for minority classes, particularly in classes 3, a substantial gap remains compared to the global model. These results indicate that ignoring vacant classes can lead to a sharp decline in accuracy for those classes, and

\*To whom correspondence should be addressed.

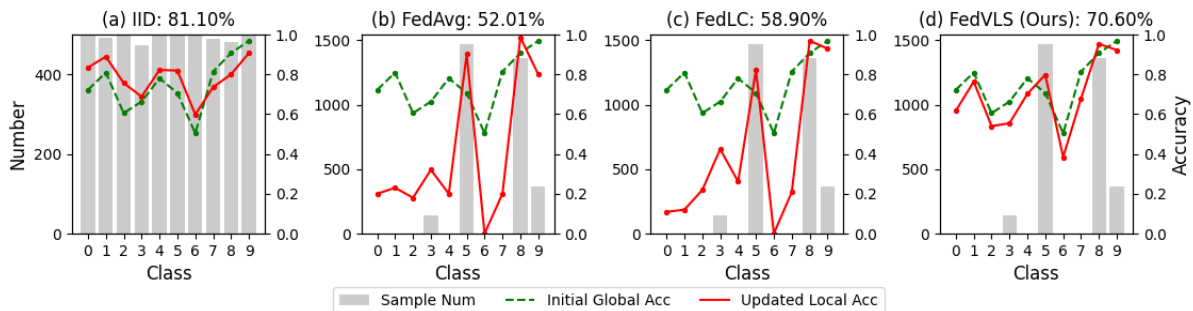


Figure 1: Class-wise accuracy of the initial global model and updated local models on IID and label-skewed CIFAR10 data distributions. (a) represents the result updating on IID local data with FedAvg (McMahan et al. 2017). (b-d) showcase the results updating on skewed data distribution with FedAvg (McMahan et al. 2017), FedLC (Zhang et al. 2022), and our FedVLS, respectively. The value (%) in each caption corresponds to the accuracy of the global model aggregated from local models.

previous methods still often misclassify minority classes.

Based on these findings, we believe improving the accuracy of vacant and minority classes is crucial for addressing the challenges posed by label skews. Therefore, we present FedVLS, a novel approach comprising two pivotal components: vacant-class distillation and logit suppression. The vacant-class distillation aims to address the performance decline related to vacant classes by distilling vital information from the global model for each client during local training. Additionally, FedVLS incorporates logit suppression, which regulates the output logit for non-label classes. This process emphasizes minimizing the predicted logit values linked to the majority class when handling minority samples, amplifying the penalty for the misclassification of minority classes. As shown in Figure 1 (d), FedVLS significantly mitigates the decline in accuracy in both vacant and minority classes of the updated local model. Consequently, FedVLS effectively reduces overfitting in client models, leading to a significant improvement in the global model’s performance. Our experiments demonstrate that FedVLS consistently outperforms current SOTA federated learning methods across various settings. Our contributions are summarized as follows:

- We find that prior federated learning methods suffer from vacant classes and propose FedVLS to distill vacant-class-aware knowledge from the global model.
- FedVLS further presents a logit suppression strategy to address the misclassification of the minority classes, thereby enhancing the generalization of local models.
- Extensive results validate the effectiveness of both components in FedVLS, outperforming previous state-of-the-art methods across diverse datasets and different degrees of label skews.

## Related Work

### Heterogeneous Federated Learning

Federated learning faces a significant challenge known as data heterogeneity, also referred to as non-identical and independently distributed (Non-IID) data (Kairouz et al. 2021; Luo et al. 2021; Shi et al. 2023b; Guo, Wang, and Geng 2024; Guo et al. 2024). This challenge encompasses issues

such as label skews and domain shifts. In this paper, we primarily focus on addressing label skews. The classic federated learning algorithm, FedAvg (McMahan et al. 2017), experiences a significant decline in performance when dealing with label skews (Li et al. 2019; Acar et al. 2021; Luo, Wang, and Wang 2024). Numerous studies have aimed to mitigate the adverse impacts of label skews. For instance, FedProx (Li et al. 2020) employs a proximal term and SCAFFOLD (Karimireddy et al. 2020) uses a variance reduction approach to constrain the update direction of local models. Additionally, MOON (Li, He, and Song 2021) and FedProc (Mu et al. 2023) utilize contrastive loss to enhance the agreement between local models and the global model. Furthermore, FedConcat (Diao, Li, and He 2024) propose model concatenation, FedMR (Fan et al. 2023) proposes a manifold reshaping approach, FedGELA (Fan et al. 2024) uses simplex Equiangular Tight to initialize the local classifier and FedGF (Lee and Yoon 2024) refine the flat minima searching to alleviate the label skews. However, these methods often fail to address the issue of vacant classes in highly skewed scenarios. We propose FedVLS to effectively mitigate the decline in class-wise accuracy of vacant classes.

### Learning from Imbalanced Data

Imbalanced data distribution is pervasive in real-world scenarios, and numerous methods have been proposed to address its impact on model performance (Cui et al. 2019; Menon et al. 2021; Tan et al. 2020; Li et al. 2022; Ma et al. 2023). Existing approaches generally fall into two categories: re-weighting (Cui et al. 2019) and logit-adjustment (Menon et al. 2021; Tan et al. 2020). However, previous works primarily discuss scenarios with long-tailed distributions (Zeng et al. 2023; Xiao et al. 2024). These methods may not be directly applicable in federated learning due to the diversity of client data distributions. In federated learning, FedLC (Zhang et al. 2022) and Calfat (Chen et al. 2022) introduce logit calibration based on the local data distribution to balance the majority and minority classes. FedLMD (Lu et al. 2023) proposes distillation masks to preserve the information of minority class. However, they often neglect vacant classes and cannot effectively handle the ac-

curacy decrease of the minority class. Our method, on the other hand, addresses both existence of vacant classes and the class imbalance between majority and minority classes, making it more practical for real-world scenarios.

### Knowledge Distillation in Federated Learning

Knowledge Distillation (KD) has been introduced to federated learning to address issues arising from variations in data distributions and model constructions across clients (Jeong et al. 2018; Itahara et al. 2021; Wu et al. 2023). FedDF (Lin et al. 2020) and FedMD (Li and Wang 2019) leverage KD to transfer the knowledge from multiple local models to the global model. However, these KD methods typically require a public dataset available to all clients on the server, which presents potential practical challenges. Recent methods, such as FEDGEN (Zhu, Hong, and Zhou 2021), DaFKD (Wang et al. 2023a), and DFRD (Luo et al. 2023), propose training a generator on the server or client to enable data-free federated knowledge distillation. However, training the generator adds computational complexity and can often be unstable in cases of extreme label skews (Wu et al. 2023). Additionally, FedNTD (Lee et al. 2022) conducts local-side distillation only for not-true labels to prevent overfitting, while FedHKD (Chen, Vikalo et al. 2023) performs local-side distillation on both logits and class prototypes to align the global and local optimization directions. However, these methods perform knowledge distillation across all classes, which may limit the retention of local models’ information about vacant classes. Our method, in contrast, applies knowledge distillation exclusively to vacant classes, preserving vital information about these categories without impacting the learning of other categories or introducing significant computational overhead.

## Method

### Preliminaries

In federated learning, we consider a scenario with  $N$  clients, where  $\mathcal{D}_i$  represents the local training data of client  $i$ . The combined data  $\mathcal{D} = \bigcup_{i=1}^N \mathcal{D}_i$  comprises the local data from all clients. These data distributions might differ across clients, encompassing situations where the local training data of some clients only contain samples from a subset of all classes. The overarching goal is to address the optimization problem as follows (McMahan et al. 2017):

$$\min_{\omega} \left[ \mathcal{L}(\omega) \stackrel{def}{=} \sum_{i=1}^N \frac{|\mathcal{D}_i|}{|\mathcal{D}|} \mathcal{L}_i(\omega) \right], \quad (1)$$

where  $\mathcal{L}_i(\omega) = \mathbb{E}_{(x,y) \sim \mathcal{D}_i} [\ell_i(f(x; \omega), y)]$  is the empirical loss of the  $i$ -th client.  $f(x; \omega)$  is the output of the model when the input  $x$  and model parameter  $\omega$  are given, and  $\ell_i$  is the loss function of the  $i$ -th client.  $|\mathcal{D}_i|$  is the number of samples on  $\mathcal{D}_i$ ,  $|\mathcal{D}|$  is the number of samples on  $\mathcal{D}$ . Here, FL expects to learn a global model that can perform well on the entire data  $\mathcal{D}$ .

### Motivation

When the local training data  $\{\mathcal{D}_i\}_{i=1}^N$  exhibit label skews, as illustrated in Figure 3 of the technical appendix, there

are variations in the quantity of data for the same category across different clients, leading the client models to excessively fit their respective local data distributions. This overfitting phenomenon causes divergence during model aggregation, subsequently resulting in inferior global performance (Yeganeh et al. 2020; Li et al. 2020; Liu et al. 2022). To address the imbalance between minority and majority classes, previous methodologies (Zhang et al. 2022; Shen, Wang, and Lv 2023) suggest calibrating logits based on the local data distribution, outlined as follows:

$$\mathcal{L}_{cal} = -\mathbb{E}_{(x,y) \sim \mathcal{D}_i} \log \left( \frac{p(y) \cdot e^{f(x; \omega)[y]}}{\sum_c p(c) \cdot e^{f(x; \omega)[c]}} \right), \quad (2)$$

where  $p(y)$  signifies the probability of class  $y$  occurring within the client’s data distribution, while  $f(x; \omega)[c]$  denotes the logit output for the  $c$ -th category. The calibration technique weights the outputs for all classes in the denominator by  $p(c)$ . However, the probability  $p(m)$  for the vacant class in the client data distribution equates to zero, leading to the weighting term for the vacant category,  $p(m) \cdot e^{f(x; \omega)[m]}$ , also becoming zero. Consequently, local models prioritize learning the majority and minority classes, gradually disregarding information associated with the vacant categories during local training. This gradual shift causes the updated direction of local models to deviate from that of the global model over time. It’s crucial to acknowledge that treating the vacant class merely as a unique minority class is insufficient, an oversight prevalent in prior methodologies. We assert this drawback significantly contributes to severe instances of local overfitting.

Our empirical observations reveal a substantial decrease in class-wise accuracy for vacant classes (such as categories 0, 1, 2, 4, 6, and 7) after the local update, as illustrated in Figure 1 (b) and (c). Notably, in specific cases (such as category 6), this accuracy even drops close to zero. The specific experimental setup and other analyses can be found in the technical appendix. By the way, we find the updated class-wise accuracy for the minority classes (e.g., category 3) continues to display a notable decline, maintaining a significant gap compared to the IID scenario. Through the analysis of the confusion matrix in Figure 2, we find that vacant and minority classes are still frequently misclassified as majority classes. Thus, we aim to develop different objectives to alleviate these two issues, respectively.

### Vacant-class Distillation

Motivated by the above observations and analyses of vacant classes, we propose to prevent the disappearance of information related to vacant classes during local training. The global model harbors valuable insights, particularly regarding the prediction of vacant classes, making it an exceptional teacher for each client. Hence, we introduce vacant-class distillation, aimed at preserving the global perspective of vacant classes for clients through knowledge distillation. To achieve this, we utilize the Kullback-Leibler Divergence loss function, as outlined below:

$$\mathcal{L}_{dis} = \mathbb{E}_{(x,y) \sim \mathcal{D}_i} \sum_{o \in \mathbb{O}} q^o(o; x) \log \left[ \frac{q(o; x)}{q^o(o; x)} \right], \quad (3)$$

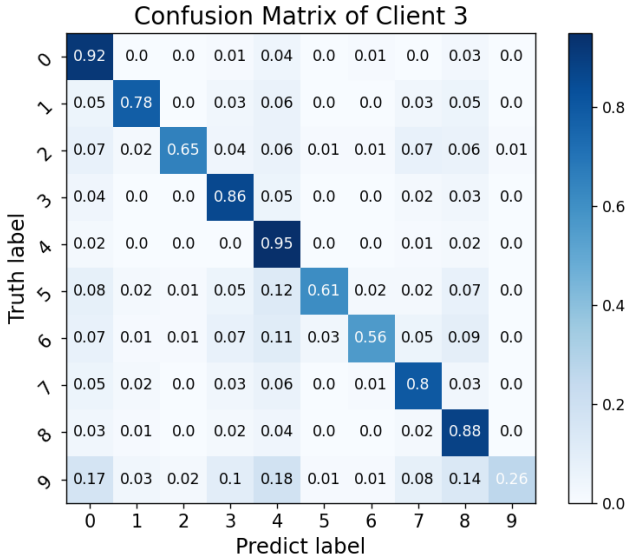


Figure 2: Confusion matrix of client 3 on CIFAR10 dataset with Dirichlet-based label skews ( $\beta = 0.5$ ) using FedLC (Zhang et al. 2022).

$$\text{where } q(o; x) = \frac{\exp(f(x; \omega)[o])}{\sum_{c \in \mathbb{O}} \exp(f(x; \omega)[c])}$$

denotes the output for the  $o$ -th class of the local model using softmax within the vacant classes, and  $q^g(o; x) = \frac{\exp(f(x; \omega^g)[o])}{\sum_{c \in \mathbb{O}} \exp(f(x; \omega^g)[c])}$  denotes the same for the global model.  $\mathbb{O}$  represents the set that contains all vacant classes within the local client data and  $\omega^g$  denotes the parameters of the global model.

Unlike FedNTD (Lee et al. 2022), which encourages the client model to closely match the global model’s output for not-true labels, thereby limiting the knowledge protection for vacant class, our loss function ensures that the local model replicates the global model’s outputs only for vacant-class labels. This approach preserves the predictive capability for vacant categories significantly. Moreover, the computational overhead introduced by this loss function is minimal, enhancing its practical implementation. Additional comparisons with other distillation-based methods and further analyses are provided in the technical appendix.

## Logit Suppression

Previous methods still often suffer from low accuracy in the minority classes of local models, a factor that requires mitigation to enhance the generalization capabilities of these models. To identify the root cause of this issue, we analyzed the confusion matrix of the local model in client 3 on the entire test dataset using FedLC (Zhang et al. 2022), where the training data distribution is shown in the fourth column of Figure 3 (a) in the technical appendix. As shown in Figure 2, minority classes (e.g., categories 2, 5, and 6) are frequently misclassified as majority classes (e.g., categories 0, 4, and 8). It is evident that, in the model’s output for minority samples, the majority class tends to have a higher logit value,

leading to the misclassification of minority classes. Therefore, we implement regularization on non-label class logits to penalize the majority class output for minority samples. To avoid non-trivial optimization over direct logits, we aim to minimize the following objective for each class:

$$\mathcal{L}_{\text{logit}}^c = \log \left( \mathbb{E}_{(x,y) \sim \mathcal{D}_i} \mathbb{I}(y \neq c) \cdot e^{f(x; \omega)[c]} \right), \quad (4)$$

where  $\mathbb{I}$  is an indicator function with value 1 when  $y \neq c$ . We use the log function to increase the proportion of loss values for minority categories. Since minority samples are prone to be more frequently misclassified into majority classes, the higher weight should be assigned to the logits of majority categories in non-labeled outputs. Therefore, we weight the loss function  $\mathcal{L}_{\text{logit}}^c$  using the probability of occurrence  $p(c)$  for each class as follows:

$$\mathcal{L}_{\text{logit}} = \sum_c p(c) \cdot \mathcal{L}_{\text{logit}}^c. \quad (5)$$

This adaptation prompts the learning process to pay more attention to the penalty for incorrectly classifying minority class samples as majority classes. As a result, the model is encouraged to refine its prediction across diverse classes, thereby improving its overall generalization capability.

## Overall Objective

As of now, we have elaborated extensively on our strategy to tackle the problem of loss of information about vacant classes in previous methods through knowledge distillation. Moreover, we mitigate the decrease in minority classes of the updated local model by regulating non-label logits directly, to further alleviate local overfitting issues. In summary, we propose the comprehensive method named FedVLS, whose objective is as follows:

$$\mathcal{L}(\omega) = \mathcal{L}_{\text{cal}} + \lambda \cdot \mathcal{L}_{\text{dis}} + \mathcal{L}_{\text{logit}}, \quad (6)$$

where  $\lambda$  is a non-negative hyperparameter to control the contribution of vacant-class distillation. In the loss function of our FedVLS, we attain new knowledge from the observed class in local data distribution using the  $\mathcal{L}_{\text{cal}}$  and  $\mathcal{L}_{\text{logit}}$ . In the meanwhile, we preserve the previous knowledge on the vacant classes by following the global model’s perspective using the  $\mathcal{L}_{\text{dis}}$ . By combining vacant-class distillation and logit suppression, FedVLS can effectively manipulate various levels of label skews. Algorithm 1 in the technical appendix shows the overflow of our method.

## Experiments

### Setups

**Datasets** We evaluate the effectiveness of our approach across various image classification datasets, including MNIST (Deng 2012), CIFAR10 (Krizhevsky 2009), CIFAR100 (Krizhevsky 2009), and TinyImageNet (Le and Yang 2015). We partitioned each dataset into distinct training and test sets. Subsequently, the training set undergoes further division into non-overlapping subsets, distributed among different clients. The global model’s performance is then assessed on the test set. We follow the settings outlined in (Li et al. 2022) and introduce two prevalent forms

Method(venue)	MNIST			CIFAR10			CIFAR100			TinyImageNet		
	$\beta = 0.5$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.5$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.5$	$\beta = 0.1$	$\beta = 0.05$	$\beta = 0.5$	$\beta = 0.1$	$\beta = 0.05$
FedAvg (AISTATS 2017)	98.96±0.00	96.69±0.00	94.77±0.44	91.46±0.55	82.00±0.75	62.90±0.95	72.22±0.34	66.18±0.35	62.13±0.09	47.02±0.40	39.90±0.27	35.21±0.47
FedProx (MLSys 2020)	98.93±0.00	96.42±0.00	94.95±0.24	92.24±0.78	82.65±1.33	63.14±0.41	72.65±0.60	66.61±0.22	62.23±0.20	45.76±0.50	40.26±0.51	35.22±0.17
MOON (CVPR 2021)	99.18±0.01	96.94±0.12	93.39±0.21	92.13±0.35	83.38±0.43	61.34±0.77	72.87±0.11	66.12±0.32	60.45±0.41	42.26±0.36	36.88±0.53	33.61±0.35
FedEXP (ICLR 2023)	97.57±0.49	91.59±0.48	92.54±1.08	92.31±0.52	83.48±1.15	63.22±0.51	72.41±0.39	66.74±0.19	62.24±0.18	47.00±0.23	40.58±0.15	34.95±0.18
FedLC (ICML 2022)	98.97±0.01	95.59±0.05	85.56±0.18	91.98±0.63	82.24±0.53	57.31±0.97	72.69±0.30	66.20±0.20	59.18±0.11	48.01±0.21	41.46±0.37	35.56±0.58
FedRS (KDD 2021)	99.03±0.00	96.67±0.01	94.60±0.40	<u>92.55±0.68</u>	<u>83.95±0.35</u>	63.17±0.57	72.99±0.20	66.84±0.25	62.19±0.06	47.95±0.43	41.77±0.25	35.82±0.20
FedSAM (ICML2022)	99.21±0.00	<u>97.24±0.00</u>	95.17±0.42	92.37±1.33	81.19±0.32	63.11±1.05	72.96±0.25	67.50±0.19	61.32±0.14	48.43±1.42	43.96±1.02	41.14±0.23
FedNTD (NeurIPS 2022)	99.15±0.04	96.67±0.17	94.30±0.71	92.46±0.19	83.23±0.22	68.71±0.27	<u>73.43±0.15</u>	68.00±0.50	63.71±0.19	48.02±1.05	45.11±0.21	40.65±0.26
FedMR (TMLR 2023)	98.95±0.02	96.73±0.08	95.34±0.50	91.98±0.55	82.09±0.42	63.54±0.69	71.94±0.36	67.57±0.37	63.75±0.24	47.21±0.53	40.35±0.26	35.94±0.46
FedLMD (MM 2023)	99.17±0.03	97.18±0.12	95.33±0.53	92.50±0.34	83.14±0.19	<u>70.50±0.29</u>	73.30±0.30	<u>68.83±0.35</u>	64.10±0.19	48.43±0.48	44.03±0.25	41.18±0.27
FedConcat (AAAI 2024)	99.04±0.01	96.99±0.11	95.02±0.47	92.45±0.29	82.83±0.21	64.30±0.28	73.27±0.28	68.57±0.34	63.74±0.13	48.45±0.44	47.32±0.21	43.44±0.21
FedGF (ICML 2024)	<u>99.22±0.00</u>	<b>97.35±0.00</b>	<u>95.36±0.28</u>	92.52±0.22	82.91±0.16	69.61±0.47	73.30±0.25	68.70±0.20	<u>64.48±0.08</u>	<u>48.52±0.23</u>	<u>47.64±0.16</u>	<u>44.71±0.20</u>
<b>FedVLS (Ours)</b>	<b>99.23±0.00</b>	<b>97.24±0.00</b>	<b>95.56±0.12</b>	<b>92.66±0.14</b>	<b>84.35±0.04</b>	<b>75.71±0.28</b>	<b>73.49±0.80</b>	<b>69.02±0.18</b>	<b>65.71±0.01</b>	<b>48.54±0.12</b>	<b>47.73±0.13</b>	<b>45.23±0.15</b>

Table 1: Performance overview for different degrees of Dirichlet-based label skews. All results are (re)produced by us and are averaged over 3 runs (mean  $\pm$  std). **Bold** is the best result, underline is the second-best.

Method(venue)	CIFAR10	CIFAR100	TinyImageNet
	$s = 2$	$s = 20$	$s = 40$
FedAvg (AISTATS 2017)	44.63±0.77	63.14±0.03	30.28±0.12
FedProx (MLSys 2020)	48.65±0.59	62.10±0.10	28.14±0.93
MOON (CVPR 2021)	38.24±1.00	57.33±0.06	26.25±0.73
FedEXP (ICLR 2023)	41.11±0.26	62.61±0.06	29.38±0.19
FedLC (ICML 2022)	55.14±0.26	61.56±0.03	26.29±1.00
FedRS (KDD 2021)	42.20±1.49	61.53±0.03	28.31±0.06
FedSAM (ICML2022)	36.97±1.18	63.50±0.01	37.55±0.10
FedNTD (NeurIPS 2022)	67.35±0.19	63.74±0.01	37.19±0.07
FedMR (TMLR 2023)	46.55±0.64	63.55±0.03	28.45±0.10
FedLMD (MM 2023)	<b>68.52±0.34</b>	63.51±0.02	32.29±0.08
FedConcat (AAAI 2024)	62.00±0.28	63.87±0.01	42.95±0.06
FedGF (ICML 2024)	66.97±0.45	<u>63.90±0.02</u>	<u>43.55±0.06</u>
<b>FedVLS (Ours)</b>	<u>68.03±0.18</u>	<b>64.95±0.01</b>	<b>43.97±0.04</b>

Table 2: Performance overview for quantity-based label skews.  $s$  presents the number of shards per client.

of label skews: Dirichlet-based and quantity-based. In the quantity-based label skews, all training data is grouped by label and allocated into shards with imbalanced quantities. The parameter  $s$  signifies the number of shards per client, regulating the level of label skews (Lee et al. 2022). In the Dirichlet-based label skews, clients receive samples for each class based on the Dirichlet distribution (Zhu et al. 2021), denoted as  $D(\beta)$ . Here, the parameter  $\beta$  controls the degree of label skews, with lower values indicating higher label skews. Notably, each client’s training data may encompass majority classes, minority classes, and even vacant classes, which is more practical.

**Models and baselines** Following a prior study (Shi et al. 2023a), our primary network architecture for all experiments, except MNIST, predominantly relies on MobileNetV2 (Sandler et al. 2018). For the MNIST, we adopt a deep neural network (DNN) containing three fully connected layers as the backbone. Our baseline models encompass conventional approaches to tackle data heterogeneity issues, including FedProx (Li et al. 2020), MOON (Li, He, and Song 2021), FedSAM (Qu et al. 2022), FedEXP (Diyansh Jhunjhunwala 2023), FedConcat (Diao, Li, and He 2024), and FedGF (Lee and Yoon 2024). To ensure a fair comparison, we also assess our method against FedRS (Li

and Zhan 2021), FedLC (Zhang et al. 2022), FedNTD (Lee et al. 2022) and FedLMD (Lu et al. 2023), which also focus on addressing label skews in federated learning.

**Implementation details** We set the number of clients  $N$  to 10 and implement full client participation. We run 100 communication rounds for all experiments on the CIFAR10/100 datasets and 50 communication rounds on the MNIST and TinyImageNet datasets. Within each communication round, local training spans 5 epochs for MNIST and 10 epochs for the other datasets. For FedConcat (Diao, Li, and He 2024) and FedGF (Lee and Yoon 2024), we followed the original paper’s settings for communication rounds and local epochs. We employ stochastic gradient descent (SGD) optimization with a learning rate of 0.01, a momentum of 0.9, and a batch size of 64. Weight decay is set to  $10^{-5}$  for MNIST and CIFAR10 and  $10^{-4}$  for CIFAR100 and TinyImageNet. The hyperparameter  $\lambda$  of FedVLS in Equation 6 is set to 0.1 for MNIST and CIFAR10, while it is set to 0.5 for CIFAR100 and TinyImageNet. Following pFedMe (T Dinh, Tran, and Nguyen 2020), we conduct three trials for each experimental setting and report the mean accuracy and standard deviation of the maximum accuracy achieved by the global model during the training process. More implementation details and experimental results can be found in the technical appendix at <https://github.com/krumpguo/FedVLS>.

## Results

### Results under various levels of label skews and datasets

Table 1 presents the performance results of various methods with different levels of Dirichlet-based label skews ( $\beta \in \{0.5, 0.1, 0.05\}$ ). Our method consistently achieves notably higher accuracy compared to other SOTA methods. As the degree of label skews increases, competing methods struggle to maintain their performance levels. For instance, FedLC (Zhang et al. 2022) experiences a substantial decline, dropping even below the performance of the classic method FedAvg when  $\beta = 0.05$ . This decline stems from each client having numerous vacant classes in extreme cases, a factor overlooked by FedLC (Zhang et al. 2022). Conversely, our method consistently upholds excellent performance, especially in highly skewed label distribution scenarios. For in-

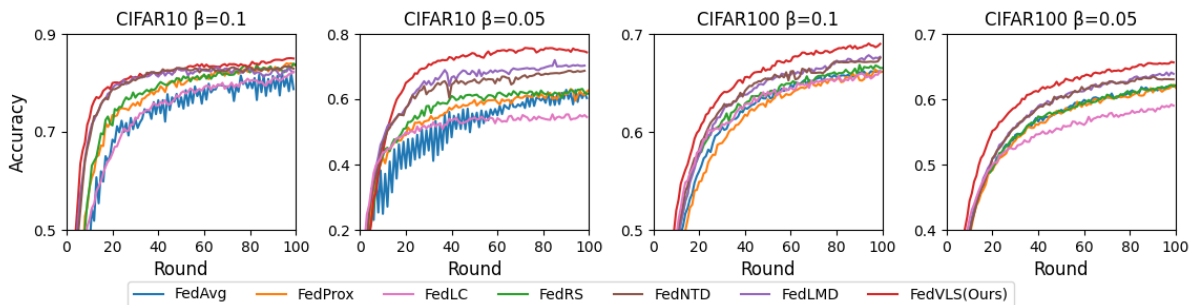


Figure 3: The test accuracy over each communication round during training for different levels of Dirichlet-based label skew ( $\beta \in \{0.1, 0.05\}$ ) on CIFAR10 and CIFAR100 datasets.

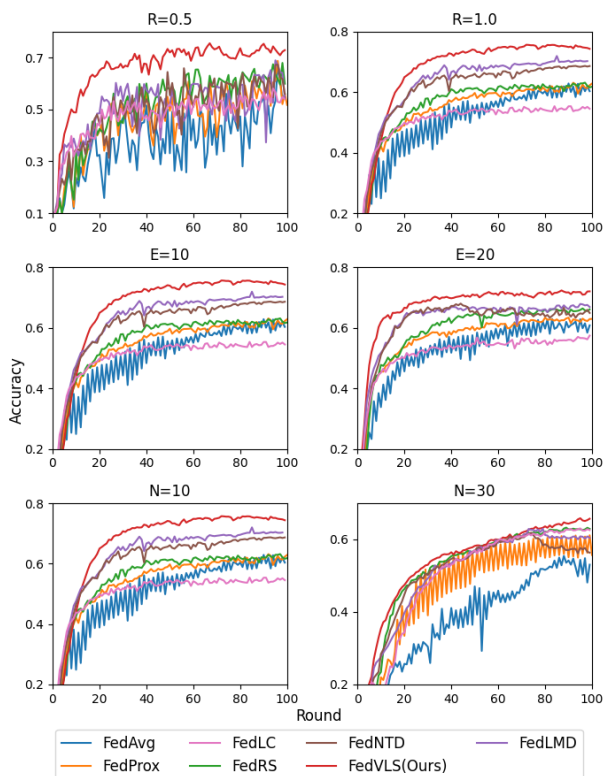


Figure 4: Sensitivity analysis on the client participating rates  $R$ , local epochs  $E$ , and client numbers  $N$ .

stance, in the case of the CIFAR10 dataset with  $\beta = 0.05$ , our method achieves an impressive test accuracy of 75.71%, surpassing FedAvg by 12.81%. This outcome highlights the efficacy of our approach in addressing the accuracy decline observed in both vacant and minority classes, effectively mitigating instances of overfitting in local data distributions. Additionally, we present the performance of these methods for quantity-based label distribution skews in Table 2, further emphasizing the superiority of our method.

**Communication efficiency** Figure 3 illustrates the accuracy over each communication round throughout the train-

Method(venue)	ResNet18		ResNet32		MobileNetV2	
	$\beta=0.1$	$\beta=0.05$	$\beta=0.1$	$\beta=0.05$	$\beta=0.1$	$\beta=0.05$
FedAvg (AISTATS 2017)	73.84	58.54	79.38	55.41	82.00	62.90
FedProx (MLSys 2020)	74.68	58.14	80.60	62.51	82.65	63.14
MOON (CVPR 2021)	74.04	55.41	76.91	51.85	83.38	61.34
FedEXP (ICLR 2023)	72.80	58.04	78.36	53.35	83.48	63.22
FedLC (ICML 2022)	73.15	48.94	77.71	55.41	82.24	57.31
FedRS (KDD 2021)	76.38	57.47	82.03	66.87	83.95	63.17
FedSAM (ICML2022)	68.42	55.42	75.66	58.88	81.19	63.11
FedNTD (NeurIPS 2022)	76.76	60.01	79.75	65.96	83.23	68.71
FedLMD (MM 2023)	77.02	65.80	81.76	68.04	83.14	70.50
FedConcat (AAAI 2024)	76.33	59.83	70.32	61.86	82.83	64.30
FedGF (ICML 2024)	76.74	64.44	81.44	67.83	82.91	69.61
<b>FedVLS (Ours)</b>	<b>78.00</b>	<b>68.33</b>	<b>82.44</b>	<b>68.84</b>	<b>84.35</b>	<b>75.71</b>

Table 3: Results of different methods under various backbones with Dirichlet-based label skews on CIFAR10 dataset.

ing process. Our method showcases quicker convergence and higher accuracy when compared to the other six methods. Due to the differences in communication rounds among FedConcat (Diao, Li, and He 2024), FedGF (Lee and Yoon 2024) and our approach, we have not included the convergence curves for these two methods. Unlike its counterparts, our approach displays a more consistent upward trend. Moreover, our method exhibits a significant improvement as the skews in the data distribution increase. These outcomes underscore the substantial communication efficiency of our method compared with other approaches.

## Analysis

**Impact of participating rates** To begin with, we analyze our model’s performance against SOTA methods across varying client participation rates. Unless specified otherwise, our experiments focus on the CIFAR10 dataset with a Dirichlet-based skew parameter of  $\beta = 0.05$ . Initially, we set the client participation rate  $R$  within the range  $\{0.5, 1.0\}$ . As illustrated in the top row of Figure 4, our method consistently outperforms other approaches across all participation rates, showcasing a faster convergence rate. Notably, as the participation rate decreases, several methods display highly unstable convergence. This instability is expected, as a lower client participation rate amplifies the divergence between randomly participating clients and the global model, resulting in erratic convergence. In contrast, our method exhibits a relatively stable convergence trend, highlighting its

$\lambda$	0.05	0.1	0.25	0.5	1
<b>CIFAR10</b>	74.70	<b>75.71</b>	75.47	75.29	74.98
<b>CIFAR100</b>	65.49	65.57	65.63	<b>65.71</b>	65.18

Table 4: Results under different values of hyperparameter  $\lambda$  with Dirichlet-based label skews ( $\beta = 0.05$ ) on CIFAR10 and CIFAR100 datasets.

$\mathcal{L}_{\text{dis}}$	$\mathcal{L}_{\text{logit}}$	<b>CIFAR10</b>	<b>CIFAR100</b>	<b>TinyImageNet</b>
$\times$	$\times$	57.31	59.18	35.56
$\times$	$\checkmark$	70.25(+12.94)	64.24(+5.06)	39.04(+3.48)
$\checkmark$	$\times$	71.53(+14.22)	65.28(+6.10)	44.90(+9.34)
$\checkmark$	$\checkmark$	75.71(+18.40)	65.71(+6.53)	45.23(+9.67)

Table 5: Effectiveness of each loss function in FedVLS with Dirichlet-based label skews ( $\beta = 0.05$ ) on various datasets. (The value) represents the improvement over the first row.

robustness to varying participation rates.

**Impact of local epochs** In this analysis, we investigate variations in the number of local epochs per communication round, represented as  $E$ , considering values from  $\{10, 20\}$ . An intriguing observation emerges, particularly noticeable when  $E$  equals 20: several methods, notably FedNTD (Lee et al. 2022), exhibit declining accuracy in the later stages of training, as depicted in the second row of Figure 4. This decline is attributed to larger  $E$  values, making these models more susceptible to overfitting local data distribution as training progresses. In contrast, our method sustains a consistent and improving performance even with larger  $E$  values and consistently outperforms all other methods.

**Impact of client numbers** To underscore the resilience of our method in scenarios involving an increasing number of clients, we divide the CIFAR10 dataset into 10 and 30 clients, showcasing their convergence curves in the final row of Figure 4. Remarkably, our method consistently outperforms the baseline methods, regardless of the number of clients. An interesting trend emerges where, with the expanding number of clients, many methods exhibit slower and less stable convergence. In contrast, FedVLS maintains a consistent trend of rapid and stable convergence across these varied client numbers. Additional ablation study results concerning participating rates, local epochs, and the number of clients can be found in the technical appendix.

**Impact of different backbones** Apart from MobileNetV2, we conduct experiments using ResNet18 and ResNet32. The skew parameter, denoted as  $\beta$ , is set to 0.1 and 0.05. The results are presented in Table 3, demonstrating our method, FedVLS, consistently outperforms the baseline methods. These experiments underscore the versatility and robustness of FedVLS in real-world federated learning scenarios employing various backbone architectures.

**Robustness to hyperparameter  $\lambda$**  To demonstrate the robustness of our method concerning hyperparameter selection, we conduct experiments using various values of  $\lambda$  on the CIFAR10 and CIFAR100 datasets. The findings, presented in Table 4, illustrate that our method

Method(venue)	<b>CIFAR10</b>	<b>CIFAR100</b>	<b>TinyImageNet</b>
FedLC (ICML 2022)	57.31	59.18	35.56
+ FedVLS (Ours)	75.71(+18.40)	65.71(+6.53)	45.23(+9.67)
FedEXP (ICLR 2023)	63.22	62.24	34.95
+ FedVLS (Ours)	75.80(+12.58)	65.94(+3.70)	44.76(+9.81)
FedSAM (ICML2022)	63.11	61.32	41.14
+ FedVLS (Ours)	75.92(+12.81)	65.46(+4.14)	48.12(+6.98)

Table 6: Results of combining FedVLS with other methods under Dirichlet-based label skews ( $\beta = 0.05$ ) across various datasets. (The values) represent the performance gains.

exhibits insensitivity to the parameter  $\lambda$ . Across  $\lambda \in \{0.05, 0.1, 0.25, 0.5, 1\}$ , our method consistently achieves approximately 75% accuracy on CIFAR10 and 65.5% on CIFAR100. This consistent performance highlights our method’s ability to deliver stable results regardless of variations in  $\lambda$  values, underscoring its robustness to hyperparameter changes.

**Effectiveness of different objectives** Our approach comprises two key objectives: vacant-classes distillation and logit suppression. The results, presented in Table 5, reveal that both vacant-classes distillation and logit suppression contribute to notable performance improvements compared to FedLC (Zhang et al. 2022). These results demonstrate the effectiveness of our two key objectives in enhancing the overall model performance in federated learning scenarios with significant label skews.

**Combination with other techniques** In this section, we integrate our method with two SOTA methods, FedEXP (Divyansh Jhunjunwala 2023) and FedSAM (Qu et al. 2022), as detailed in Table 6. The combination of our method with FedEXP (Divyansh Jhunjunwala 2023) and FedSAM (Qu et al. 2022) results in improved performance. This enhancement is reasonable because FedEXP focuses on optimizing the server update for an improved learning rate, and FedSAM emphasizes local gradient descent to achieve a smoother loss landscape. These elements complement well with our core idea, which finally results in enhanced performance when combined.

## Conclusion

We have observed that existing federated learning methods always perform poorly in vacant and minority classes, under skewed label distribution across clients. To overcome these challenges, we introduce FedVLS—an innovative methodology integrating vacant-class distillation and logit suppression simultaneously. The vacant-class distillation extracts pertinent knowledge regarding vacant classes from the global model for each client, while logit suppression is implemented to directly regularize non-label class logits, addressing the imbalance among majority and minority classes. Extensive results affirm the effectiveness of both components, surpassing previous state-of-the-art methods across diverse datasets and varying degrees of label skews. In future work, we will conduct a theoretical analysis of FedVLS, 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) and the Young Elite Scientists Sponsorship Program by CAST (2023QNRC001).

## References

- Acar, D. A. E.; Zhao, Y.; Navarro, R. M.; Mattina, M.; Whatmough, P. N.; and Saligrama, V. 2021. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263*.
- Chai, D.; Wang, L.; Yang, L.; Zhang, J.; Chen, K.; and Yang, Q. 2024. A survey for federated learning evaluations: Goals and measures. *IEEE Transactions on Knowledge and Data Engineering*, 36(10): 5007–5024.
- Chen, C.; Liu, Y.; Ma, X.; and Lyu, L. 2022. Calfat: Calibrated federated adversarial training with label skewness. In *Proc. NeurIPS*.
- Chen, H.; Vikalo, H.; et al. 2023. The Best of Both Worlds: Accurate Global and Personalized Models through Federated Learning with Data-Free Hyper-Knowledge Distillation. In *Proc. ICLR*.
- Cui, Y.; Jia, M.; Lin, T.-Y.; Song, Y.; and Belongie, S. 2019. Class-balanced loss based on effective number of samples. In *Proc. CVPR*.
- Deng, L. 2012. The mnist database of handwritten digit images for machine learning research [best of the web]. *IEEE Signal Processing Magazine*, 29(6): 141–142.
- Diao, Y.; Li, Q.; and He, B. 2024. Exploiting Label Skews in Federated Learning with Model Concatenation. In *Proc. AAAI*.
- Divyansh Jhunjhunwala, G. J., Shiqiang Wang. 2023. Fed-ExP: Speeding up Federated Averaging Via Extrapolation. In *Proc. ICLR*.
- Fan, Z.; Yao, J.; Han, B.; Zhang, Y.; Wang, Y.; et al. 2024. Federated Learning with Bilateral Curation for Partially Class-Disjoint Data. *Proc. NeurIPS*.
- Fan, Z.; Yao, J.; Zhang, R.; Lyu, L.; Wang, Y.; and Zhang, Y. 2023. Federated Learning under Partially Disjoint Data via Manifold Reshaping. *Transactions on Machine Learning Research*, 2023.
- Guo, S.; Wang, H.; and Geng, X. 2024. Dynamic heterogeneous federated learning with multi-level prototypes. *Pattern Recognition*, 153: 110542.
- Guo, S.; Wang, H.; Lin, S.; Kou, Z.; and Geng, X. 2024. Addressing Skewed Heterogeneity via Federated Prototype Rectification With Personalization. *IEEE Transactions on Neural Networks and Learning Systems*, 1–13.
- Hsu, T.-M. H.; Qi, H.; and Brown, M. 2019. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*.
- Itahara, S.; Nishio, T.; Koda, Y.; Morikura, M.; and Yamamoto, K. 2021. Distillation-based semi-supervised federated learning for communication-efficient collaborative training with non-iid private data. *IEEE Transactions on Mobile Computing*, 22(1): 191–205.
- Jeong, E.; Oh, S.; Kim, H.; Park, J.; Bennis, M.; and Kim, S.-L. 2018. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. In *Proc. NeurIPS Workshops*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *Proc. ICML*.
- Konečný, J.; McMahan, H. B.; Ramage, D.; and Richtárik, P. 2016. Federated optimization: Distributed machine learning for on-device intelligence. *arXiv preprint arXiv:1610.02527*.
- Krizhevsky, A. 2009. Learning Multiple Layers of Features from Tiny Images. *Master's thesis, University of Tront*.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7): 3.
- Lee, G.; Jeong, M.; Shin, Y.; Bae, S.; and Yun, S.-Y. 2022. Preservation of the global knowledge by not-true distillation in federated learning. In *Proc. NeurIPS*.
- Lee, T.; and Yoon, S. W. 2024. Rethinking the Flat Minima Searching in Federated Learning. In *Proc. ICML*.
- Li, D.; and Wang, J. 2019. Fedmd: Heterogenous federated learning via model distillation. In *Proc. NeurIPS Workshops*.
- Li, Q.; Diao, Y.; Chen, Q.; and He, B. 2022. Federated learning on non-iid data silos: An experimental study. In *Proc. ICDE*.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proc. CVPR*.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3): 50–60.
- Li, X.; Huang, K.; Yang, W.; Wang, S.; and Zhang, Z. 2019. On the convergence of fedavg on non-iid data. In *Proc. ICLR*.
- Li, X.-C.; and Zhan, D.-C. 2021. Fedrs: Federated learning with restricted softmax for label distribution non-iid data. In *Proc. KDD*.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble distillation for robust model fusion in federated learning. In *Proc. NeurIPS*.
- Liu, D.; Bai, L.; Yu, T.; and Zhang, A. 2022. Towards Method of Horizontal Federated Learning: A Survey. In *Proc. BigDIA*.
- Lu, J.; Li, S.; Bao, K.; Wang, P.; Qian, Z.; and Ge, S. 2023. Federated Learning with Label-Masking Distillation. In *Proc. ACM-MM*.
- Luo, K.; Wang, S.; Fu, Y.; Li, X.; Lan, Y.; and Gao, M. 2023. DFRD: Data-Free Robustness Distillation for Heterogeneous Federated Learning. In *Proc. NeurIPS*.
- Luo, M.; Chen, F.; Hu, D.; Zhang, Y.; Liang, J.; and Feng, J. 2021. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. In *Proc. NeurIPS*.

- Luo, Z.; Wang, Y.; and Wang, Z. 2024. Federated Local Compact Representation Communication: Framework and Application. *Machine Intelligence Research*, 21(6): 1103–1120.
- Luo, Z.; Wang, Y.; Wang, Z.; Sun, Z.; and Tan, T. 2022. Disentangled federated learning for tackling attributes skew via invariant aggregation and diversity transferring. *arXiv preprint arXiv:2206.06818*.
- Ma, Y.; Jiao, L.; Liu, F.; Yang, S.; Liu, X.; and Li, L. 2023. Curvature-Balanced Feature Manifold Learning for Long-Tailed Classification. In *Proc. CVPR*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Proc. International Conference on Artificial Intelligence and Statistics*.
- Menon, A. K.; Jayasumana, S.; Rawat, A. S.; Jain, H.; Veit, A.; and Kumar, S. 2021. Long-tail learning via logit adjustment. In *Proc. ICLR*.
- Mu, X.; Shen, Y.; Cheng, K.; Geng, X.; Fu, J.; Zhang, T.; and Zhang, Z. 2023. Fedproc: Prototypical contrastive federated learning on non-iid data. *Future Generation Computer Systems*, 143: 93–104.
- Qu, Z.; Li, X.; Duan, R.; Liu, Y.; Tang, B.; and Lu, Z. 2022. Generalized federated learning via sharpness aware minimization. In *Proc. ICML*.
- Reguieg, H.; El Hanjri, M.; El Kamili, M.; and Kobbane, A. 2023. A Comparative Evaluation of FedAvg and Per-FedAvg Algorithms for Dirichlet Distributed Heterogeneous Data. In *Proc. International Conference on Wireless Networks and Mobile Communications*.
- Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. CVPR*.
- Sheller, M. J.; Edwards, B.; Reina, G. A.; Martin, J.; Pati, S.; Kotrotsou, A.; Milchenko, M.; Xu, W.; Marcus, D.; Colen, R. R.; et al. 2020. Federated learning in medicine: facilitating multi-institutional collaborations without sharing patient data. *Scientific reports*, 10(1): 12598.
- Shen, Y.; Wang, H.; and Lv, H. 2023. Federated Learning with Classifier Shift for Class Imbalance. *arXiv preprint arXiv:2304.04972*.
- Shi, Y.; Liang, J.; Zhang, W.; Tan, V. Y.; and Bai, S. 2023a. Towards Understanding and Mitigating Dimensional Collapse in Heterogeneous Federated Learning. In *Proc. ICLR*.
- Shi, Y.; Liang, J.; Zhang, W.; Xue, C.; Tan, V. Y.; and Bai, S. 2023b. Understanding and Mitigating Dimensional Collapse in Federated Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5): 2936–2949.
- T Dinh, C.; Tran, N.; and Nguyen, J. 2020. Personalized federated learning with moreau envelopes. In *Proc. NeurIPS*.
- Tan, J.; Wang, C.; Li, B.; Li, Q.; Ouyang, W.; Yin, C.; and Yan, J. 2020. Equalization loss for long-tailed object recognition. In *Proc. CVPR*.
- Wang, H.; Li, Y.; Xu, W.; Li, R.; Zhan, Y.; and Zeng, Z. 2023a. DaFKD: Domain-aware Federated Knowledge Distillation. In *Proc. CVPR*.
- Wang, Y.; Li, R.; Tan, H.; Jiang, X.; Sun, S.; Liu, M.; Gao, B.; and Wu, Z. 2023b. Federated Skewed Label Learning with Logits Fusion. *arXiv preprint arXiv:2311.08202*.
- Weyand, T.; Araujo, A.; Cao, B.; and Sim, J. 2020. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *Proc. CVPR*.
- Wu, Z.; Sun, S.; Wang, Y.; Liu, M.; Jiang, X.; and Li, R. 2023. Survey of knowledge distillation in federated edge learning. *arXiv preprint arXiv:2301.05849*.
- Xiao, Z.; Chen, Z.; Liu, S.; Wang, H.; Feng, Y.; Hao, J.; Zhou, J. T.; Wu, J.; Yang, H.; and Liu, Z. 2024. Fed-grab: Federated long-tailed learning with self-adjusting gradient balancer. In *Proc. NeurIPS*.
- Yang, H.; Fang, M.; and Liu, J. 2021. Achieving linear speedup with partial worker participation in non-iid federated learning. In *Proc. ICLR*.
- Ye, M.; Fang, X.; Du, B.; Yuen, P. C.; and Tao, D. 2023. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3): 79:1–79:44.
- Yeganeh, Y.; Farshad, A.; Navab, N.; and Albarqouni, S. 2020. Inverse distance aggregation for federated learning with non-iid data. In *Proc. MICCAI Workshops*.
- Zeng, Y.; Liu, L.; Liu, L.; Shen, L.; Liu, S.; and Wu, B. 2023. Global Balanced Experts for Federated Long-Tailed Learning. In *Proc. ICCV*.
- Zhang, J.; Li, C.; Qi, J.; and He, J. 2023. A Survey on Class Imbalance in Federated Learning. *arXiv preprint arXiv:2303.11673*.
- Zhang, J.; Li, Z.; Li, B.; Xu, J.; Wu, S.; Ding, S.; and Wu, C. 2022. Federated learning with label distribution skew via logits calibration. In *Proc. ICML*.
- Zhu, H.; Xu, J.; Liu, S.; and Jin, Y. 2021. Federated learning on non-IID data: A survey. *Neurocomputing*, 465: 371–390.
- Zhu, Z.; Hong, J.; and Zhou, J. 2021. Data-free knowledge distillation for heterogeneous federated learning. In *Proc. ICML*.