

# FedVCK: Non-IID Robust and Communication-Efficient Federated Learning via Valuable Condensed Knowledge for Medical Image Analysis

Guochen Yan<sup>1,3,4</sup>, Luyuan Xie<sup>2,3,4</sup>, Xinyi Gao<sup>5</sup>, Wentao Zhang<sup>6</sup>,  
Qingni Shen<sup>2,3,4\*</sup>, Yuejian Fang<sup>2,3,4</sup>, Zhonghai Wu<sup>2,3,4†</sup>

<sup>1</sup>School of Computer Science, Peking University, Beijing, China

<sup>2</sup>School of Software and Microelectronics, Peking University, Beijing, China

<sup>3</sup>PKU-OCTA Laboratory for Blockchain and Privacy Computing, Peking University, Beijing, China

<sup>4</sup>National Engineering Research Center for Software Engineering, Peking University, Beijing, China

<sup>5</sup>The University of Queensland, Brisbane, Australia

<sup>6</sup>Center for Machine Learning Research, Peking University, Beijing, China

Guochen Yan@outlook.com, 2201110745@stu.pku.edu.cn, xinyi.gao@uq.edu.au, wentao.zhang@pku.edu.cn,  
qingnishen@ss.pku.edu.cn, fangyj@ss.pku.edu.cn, wuzh@pku.edu.cn

## Abstract

Federated learning has become a promising solution for collaboration among medical institutions. However, data owned by each institution would be highly heterogeneous and the distribution is always non-independent and identical distribution (non-IID), resulting in client drift and unsatisfactory performance. Despite existing federated learning methods attempting to solve the non-IID problems, they still show marginal advantages but rely on frequent communication which would incur high costs and privacy concerns. In this paper, we propose a novel federated learning method: **Federated learning via Valuable Condensed Knowledge (FedVCK)**. We enhance the quality of condensed knowledge and select the most necessary knowledge guided by models, to tackle the non-IID problem within limited communication budgets effectively. Specifically, on the client side, we condense the knowledge of each client into a small dataset and further enhance the condensation procedure with latent distribution constraints, facilitating the effective capture of high-quality knowledge. During each round, we specifically target and condense knowledge that has not been assimilated by the current model, thereby preventing unnecessary repetition of homogeneous knowledge and minimizing the frequency of communications required. On the server side, we propose relational supervised contrastive learning to provide more supervision signals to aid the global model updating. Comprehensive experiments across various medical tasks show that FedVCK can outperform state-of-the-art methods, demonstrating that it's non-IID robust and communication-efficient.

## Introduction

Federated learning has become increasingly attractive since it allows collaborative training among sensitive institutions without direct data sharing. However, in reality, each medical institution would have its specialization, and the private

Method	Syn. Data Quality	Knowledge Selection
FedGen	calibrate classifiers	heuristic
FedMix	distorted sample	repeated
FedGAN	match single sample	repeated
DFRD	infidelity	heuristic
FedDM	only match final feature	repeated
DESA	only match final feature	repeated
FedVCK	<b>latent dist. constraints</b>	<b>model-guided selection</b>

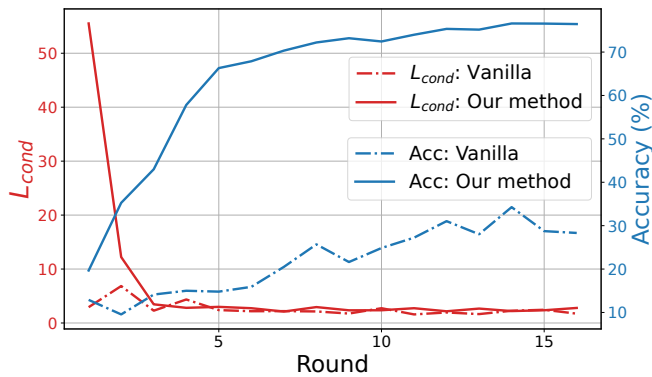
Table 1: Representative data-centric methods' problems of 1) synthesis data quality and 2) knowledge selection in the synthesis. 'heuristic' indicates they adopt heuristic diversity loss with no relation to the need of models. 'repeated' indicates they do nothing and thus select data with repeated knowledge in synthesis. In contrast, we adopt latent distribution constraints and model-guided selection respectively.

data are highly related to the regional demographic characteristics. The data owned by each client are non-independent and identical distribution (non-IID), exhibiting significant data heterogeneity and imbalance. Under this scenario, federated learning methods suffer a global model with unsatisfactory performance due to the model divergence and client drift phenomenon (Li et al. 2019, 2022a). Meanwhile, frequent communication between heterogeneous and dispersed institutions would incur high communication costs, delays, and complex administrative procedures, with increasing privacy and safety risks (Zhu, Liu, and Han 2019; Mothukuri et al. 2021). A non-IID robust and communication-efficient federated learning method is desired.

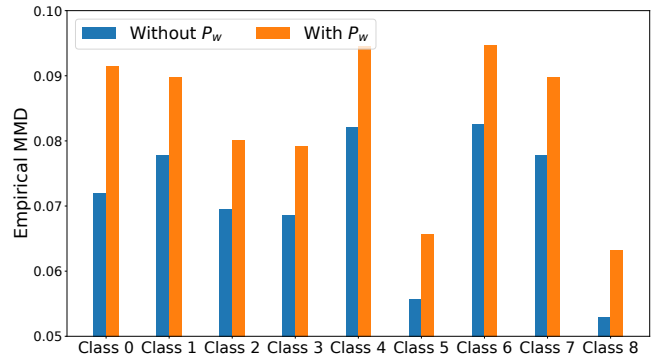
Many federated learning methods are proposed to cope with non-IID problems by modifying local training process (Li et al. 2020; Li, He, and Song 2021; Zhou, Zhang, and Tsang 2023) or global aggregation process (Chen and Chao 2020; Lin et al. 2020; Zheng et al. 2023). However, they are mainly **model-centric**. They focus on mitigating model parameter-level divergence indirectly under a typical paradigm of local training and global aggregation, leading to marginal advantages in performance and communication

\*Corresponding author: qingnishen@ss.pku.edu.cn

†Corresponding author: wuzh@pku.edu.cn



(a) We use the distribution matching-based method to condense knowledge on the OrganC dataset. Without our latent distribution constraints (Vanilla, dashed line),  $L_{cond}$  would be easily reduced in each round but the model’s performance struggles to improve with the condensed knowledge, demonstrating the low-quality problem of vanilla methods.



(b) We measure the average MMD of condensed knowledge classwisely between adjacent rounds on the Path dataset. Greater MMD indicates a larger distribution difference. The vanilla selection causes more knowledge repetition between rounds. Our model-guided selection ( $P_w$ ) ensures that the condensed knowledge between adjacent rounds exhibits greater differences.

Figure 1: Illustration of the low synthesized data quality problem in Figure (a) and repeated knowledge problem in Figure (b).

costs under severe non-IID scenarios (Li et al. 2022a).

The model divergence originates from the data divergence (Zhao et al. 2018), thus mitigating the data-level divergence would be more essential to tackle the non-IID problems. Recently, various **data-centric** federated learning methods attempted to share virtual or synthesized data to mitigate data divergence. They synthesize diverse objectives including latent features (Zhu, Hong, and Zhou 2021), approximated real data (Li et al. 2022b; Zhu and Luo 2022; Yoon et al. 2021), inverted data (Zhang et al. 2022; Wang et al. 2024a), condensed data (Xiong et al. 2023; Huang et al. 2024; Wang et al. 2024b) and so on. However, under severe non-IID scenarios, these methods still face problems because of: 1) **low synthesized data quality**. For instance, mix-up would distort data (Verma et al. 2019). The inverted data is of infidelity with biased models. And the advanced dataset condensation cannot effectively extract subtle and meaningful knowledge which we demonstrate in Figure 1a. These low-quality data would fail to guide the model training; 2) **repeated knowledge**. The data are randomly selected to synthesize virtual data, and their value and importance to the current model are not considered. Thus, the knowledge contained tends to be homogeneous and unnecessarily repeated (see Figure 1b), thus cannot effectively update the model after several rounds. We summarize problems in the above two aspects of representative data-centric federated learning methods in Table 1. Additionally, some synthesis methods would incur privacy concerns, and most methods are **not communication-efficient**. They still face challenges to achieve a satisfactory performance under limited communication rounds in non-IID scenarios.

Motivated by the above limitations, we propose a novel data-centric **F**ederated learning method via **V**aluable **C**ondensed **K**nowledge (FedVCK). Our method includes two parts, valuable knowledge condensation on the client side and relational supervised learning-aided updating on the

server side. Specifically, we condense each client’s knowledge into a small dataset. To ensure condensing high-quality knowledge, we propose latent distribution constraints to better capture subtle and meaningful knowledge in latent spaces. To minimize redundancy in each round of knowledge condensation, we explicitly measure the missing knowledge of the current model and select the most necessary knowledge in condensation on each client. On the server side, we identify the hard negative classes for each class and propose a relational supervised contrastive learning to enhance the supervision signals during model updating. Due to the balanced, high-quality, unrepeatable, and necessary condensed knowledge, the training of the global model is insulated from the effects of non-IID problems and can achieve enhanced performance within limited communication rounds (e.g. 10). Moreover, our method only condenses task-related high-level knowledge with random noise initialization, thereby facilitating privacy protection. Our main contributions are summarized as follows:

- We propose a novel data-centric federated learning method: FedVCK, for collaborative medical image analysis. FedVCK is robust to severe non-IID scenarios and communication efficient with valuable knowledge.
- On the client side, we propose model-guided selection to sample the most needed knowledge each round to avoid unnecessary repetition. We also propose latent distribution constraints to enhance the quality of knowledge.
- On the server side, we identify the hard negative classes and propose relational supervised contrastive learning to enhance supervised learning in model updating.
- We conduct comprehensive experiments and results show that our method achieves better predictive performance, especially under limited communication budgets. We also conduct experiments to verify the privacy-preserving ability and generality of our method.

## Related Works

The data owned by each client is typically highly heterogeneous and does not follow an independent and identical distribution. Under severe non-IID scenarios, models trained on clients tend to be highly biased and divergent, a phenomenon known as client drift. Aggregating these biased and divergent client models at the server often results in suboptimal performance. Many model-centric methods focus on modifying local training process, such as introducing regularization or contrastive terms to reduce divergence (Li et al. 2020; Acar et al. 2021; Li, He, and Song 2021; Xie et al. 2024c,b,a) or improving aggregation process (Lin et al. 2020; Chen and Chao 2020; Zheng et al. 2023). They try to alleviate the client drift from the model parameter level.

The model divergence originates from the data divergence (Zhao et al. 2018). Directly reducing the difference in data distribution would reduce the model divergence fundamentally. Recently, data-centric federated learning methods have drawn attention since they can synthesize and then share virtual synthesized data to mitigate the non-IID problem in a data-centric manner. Besides that FedGen (Zhu, Hong, and Zhou 2021) which generates virtual representation, various format data are synthesized on the server or clients. FedMix (Yoon et al. 2021) broadcasts the mix-up data to approximate the real data. FedGAN (Nguyen et al. 2021) and SDA-FL (Li et al. 2022b) train and share GANs to imitate real data to support COVID-19 detection. FedFTG (Zhang et al. 2022) and advanced version DFRD (Wang et al. 2024a) use model inversion (Yin et al. 2020) to generate data for knowledge distillation. FedDM (Xiong et al. 2023) condenses the knowledge to update the global model. DESA (Huang et al. 2024) distills anchor data and broadcasts them to enable mutual regularization and distillation among clients.

## Proposed Method

### Overview

The overview of FedVCK is shown in Figure 2. In short, it consists of two parts: valuable knowledge condensation on the client side and relational supervised learning-aided updating on the server side. On the client side, we borrow distribution matching techniques in dataset condensation and optimize the learnable dataset to condense knowledge from local data. To ensure quality, we record dynamic distribution statistics of the local data batch in each encoder layer and replace the statistics during embedding learnable knowledge as fixed constraints, which could force the latent distribution of the condensed knowledge to capture subtle and meaningful knowledge of different levels. To minimize redundancy in each round of condensation, we explicitly measure the prediction error on each sample and select the data on which the model performs poorly. We consider such data critical as it contains knowledge not yet captured by the current model. By focusing more on these important samples, the condensation process ensures that the condensed knowledge complements the global model’s missing capabilities. On the server side, we collect the condensed knowledge dataset and train the global model with supervised learning and relational

contrastive learning. We first identify hard negative classes for each class where the global model tends to mispredict by uploaded logit prototypes. Then we use supervised contrastive learning in a bootstrap manner to draw the features of the same class closer to their prototypes and push the features away from their hard negative classes’ prototypes. We will introduce our designs in detail in the following sections.

### Preliminary: Dataset Condensation

The objective of dataset condensation (Wang et al. 2018; Yu, Liu, and Wang 2023; Gao et al. 2024a,b) is to condense knowledge from a large dataset into a small learnable dataset, which could be used to train models to achieve comparable performance. Distribution matching is an advanced method widely used in dataset condensation.

**Distribution matching.** The intuition behind is to optimize a small dataset  $\mathcal{S}$  to match the latent feature distribution of local data  $\mathcal{T}$  by minimizing the distance to the latent features of local data with maximum mean discrepancy (MMD) (Gretton et al. 2012; Zhao and Bilén 2023):

$$\arg \min_{\mathcal{S}} \sup_{\|\psi_{\theta}\|_{\mathcal{H}} \leq 1} (\mathbb{E}[\psi_{\theta}(\mathcal{T})] - \mathbb{E}[\psi_{\theta}(\mathcal{S})]), \quad (1)$$

where  $\mathcal{H}$  is reproducing kernel Hilbert space (RKHS),  $\psi_{\theta}$  is the shared embedding function to map the input to its latent feature, parameterized by a multi-layer encoder. In practice, We minimize the estimated empirical MMD loss by classwisely align the latent feature distributions to optimize  $\mathcal{S}$ :

$$L_{cond} = \sum_{c=0}^{C-1} \left\| \frac{1}{|B_c|} \sum_{x_i \in B_c} \psi_{\theta}(x_i) - \frac{1}{|\mathcal{S}_c|} \sum_{s_i \in \mathcal{S}_c} \psi_{\theta}(s_i) \right\|^2, \quad (2)$$

where  $C$  is the number of classes,  $\mathcal{T}_c$  is the local data with class  $c$ ,  $B_c$  is a batch randomly sampled from  $\mathcal{T}_c$  with a uniform distribution, and  $\mathcal{S}_c$  is the knowledge dataset corresponding to class  $c$ . To enable the high-order estimation, we choose to align the latent feature distributions in an RKHS with kernel  $\mathcal{K}$  (Zhang et al. 2024), and minimize the following empirical MMD as condensation loss:

$$L_{cond} = \sum_{c=0}^{C-1} \sum_{B_c \overset{P_c}{\sim} \mathcal{T}_c} \hat{\mathcal{K}}_{B_c, B_c} + \hat{\mathcal{K}}_{\mathcal{S}_c, \mathcal{S}_c} - 2\hat{\mathcal{K}}_{B_c, \mathcal{S}_c}, \quad (3)$$

where  $\hat{\mathcal{K}}_{X, Y} = \frac{1}{|X| \cdot |Y|} \sum_{i=1}^{|X|} \sum_{j=1}^{|Y|} \mathcal{K}(\psi_{\theta}(x_i), \psi_{\theta}(y_j))$ ,  $\{x_i\}_{i=1}^{|X|} \sim X$ ,  $\{y_j\}_{j=1}^{|Y|} \sim Y$ . The kernel function  $\mathcal{K}$  can be a linear kernel, inner-product kernel, or Gaussian kernel.

**Knowledge initialization.** There are several manners to initialize learnable knowledge dataset  $\mathcal{S}$  whose format is the same as real data. To best prevent the privacy leak of the local data, we choose random Gaussian noise  $\mathcal{N}(0, 1)$  to initialize the knowledge dataset, making the knowledge can only be condensed by matching the latent distributions. The condensed knowledge dataset  $\mathcal{S}$  would contain no individual and privacy information in pixel space and the adversary can hardly infer the membership from the condensed knowledge datasets (Dong, Zhao, and Lyu 2022).

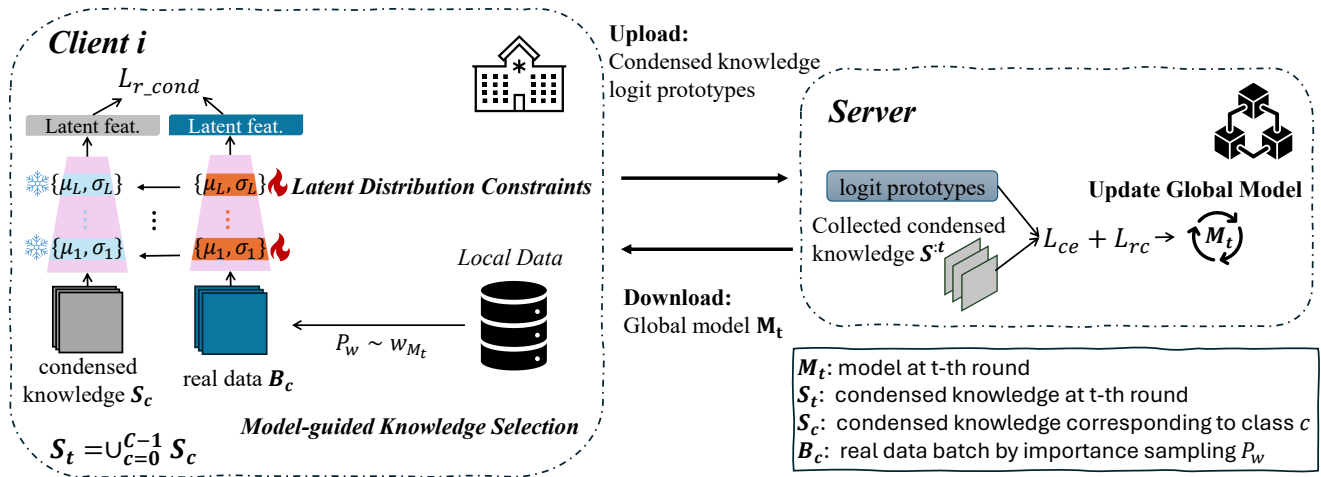


Figure 2: Overview of FedVCK. **On the client side**, we sample local data by importance sampling guided by the current model and then impose latent distribution constraints in optimization. We upload the condensed knowledge dataset and logit prototypes to the server. **On the server side**, we use cross entropy loss and relational contrastive loss to update the global model.

### Latent Distribution Constraints

By optimizing  $\mathcal{S}$  with  $L_{cond}$  in Eq. 2 or Eq. 3, the condensed knowledge dataset  $\mathcal{S}$  can replace the real data to effectively train models. However, it's challenging to ensure the condensation quality when condensing the knowledge from local data into a small random noise-initialized dataset, because random noises have no prior about the local data. Moreover, the condensation loss ( $L_{cond}$ ) is not sufficient to guide the learning of subtle meaningful knowledge since the latent features of  $\mathcal{S}$  can take shortcuts to over-fit the latent features of local data. To assess the deficiency of the vanilla optimization process, we show in Figure 1a that while  $L_{cond}$  is easily reduced at each round, it can not effectively condense meaningful knowledge to improve performance effectively. Additionally, matching the distribution of the final representation of the encoder neglects the previous intermediate latent feature distributions.

To enhance the effectiveness and ensure consistent distribution of all latent features, we transfer dynamic distribution statistics (mean and variation) from the local real data to the condensed knowledge data during the condensation procedure. Compared to (Yin et al. 2020), our approach does not require the addition of an extra loss term, enabling a more flexible and efficient condensation procedure. Specifically, we first record the distribution statistics of each layer  $\{\{\mu_1, \sigma_1\}, \dots, \{\mu_L, \sigma_L\}\}$  with a  $L$  layer encoder when embedding a batch of local data. Then the statistics during embedding condensed data are replaced and fixed with recorded statistics:

$$\begin{aligned}
 s_i^{(1)} &= Norm(\psi_1(s_i), \{\mu_1, \sigma_1\}), \\
 s_i^{(2)} &= Norm(\psi_2(s_i^{(1)}), \{\mu_2, \sigma_2\}), \\
 &\dots \\
 \psi_\theta(s_i) &= s_i^{(L)} = Norm(\psi_L(s_i^{(L-1)}), \{\mu_L, \sigma_L\}),
 \end{aligned} \tag{4}$$

where  $\psi_l$  is  $l$ -th layer of encoder  $\psi_\theta$  and  $Norm$  denotes

batch normalization (Ioffe and Szegedy 2015). The distribution statistics constraint could force the optimization process to consider intermediate latent distributions and prevent it from taking shortcuts. Thus the quality of condensed knowledge can be largely improved.

### Model-guided Knowledge Selection

If we uniformly sample real data from local data to conduct condensation, the condensed knowledge in each round will be repeated and homogeneous, dominated by simple and easy-to-learn knowledge. It would be less beneficial to further improve the performance of the global model. However, from the model perspective, we find that the importance of the knowledge contained in each sample varies. The current global model may perform well on some local data but lacks the ability to make correct predictions on others, exposing that the current model lacks some knowledge. The data containing the missing knowledge would be more important at this round and it's better to focus on condensing knowledge from these data to complement the model knowledge. Specifically, we first measure the importance of each data sample explicitly by model prediction error as  $t$ -th round:

$$w_{M_t}(x_i) = \frac{1}{1 + e^{-err_t(x_i)+b}}, \tag{5}$$

where the  $err_t(x_i)$  refers to the prediction error on  $x_i$  of current model  $M_t$  at  $t$ -th round and  $b$  is a constant to control the scale range. The higher the prediction error, the more important  $x_i$  would be. Here we adopt the cross-entropy loss as the prediction error with the current model  $M_t$ :

$$err_t(x_i) = L_{ce}(y_i, M_t(x_i)). \tag{6}$$

The current model is usually not well-trained and may over-fit on limited uploaded condensed knowledge, the distribution of loss would be skewed and less calibrated to reflect the proper importance relation. We propose the self-ensemble of the current model  $M_t$  and previous model

$\mathbf{M}_{t-1}$  to smooth and regularize the distribution of loss. Thus the desired knowledge can be condensed progressively. we refine the prediction error of Eq. 6 as:

$$err_t(x_i) = L_{ce}(y_i, \alpha \mathbf{M}_t(x_i) + (1 - \alpha) \mathbf{M}_{t-1}(x_i)), \quad (7)$$

where  $\alpha$  is a hyper-parameter. With refined prediction error, we can calculate the refined importance in Eq. 5 and replace the uniform sampling  $P_u$  in Eq. 3 with importance sampling  $P_w$  conditioned on model  $\mathbf{M}_t$  and  $\mathbf{M}_{t-1}$ :

$$P_w(x_i | \mathbf{M}_t, \mathbf{M}_{t-1}) = \frac{w_{\mathbf{M}_t}(x_i)}{\sum_{x_j} w_{\mathbf{M}_t}(x_j)}. \quad (8)$$

We then refine the condensation loss of Eq. 3 as :

$$L_{r.cond} = \sum_{c=0}^{C-1} \hat{\mathcal{K}}_{B_c^{P_w}, B_c^{P_w}} + \hat{\mathcal{K}}_{S_c, S_c} - 2\hat{\mathcal{K}}_{B_c^{P_w}, S_c}, \quad (9)$$

where data in each batch  $B_c^{P_w}$  is sampled based on  $P_w$ . Thus the condensation process in the  $t$ -th round can be regarded as a biased variant of Eq. 1 where the where the expectation over  $\mathcal{T}$  is replaced by a weighted expectation under  $P_w$ :

$$\arg \min_S \sup_{\|\psi_\theta\|_{\mathcal{H}} \leq 1} (\mathbb{E}_{P_w}[\psi_\theta(\mathcal{T})] - \mathbb{E}[\psi_\theta(S)]), \quad (10)$$

Note that the current model  $\mathbf{M}_t$  is dynamically updating. We measure the importance and derive the importance sampling  $P_w$  at each round. Thus, the condensed knowledge in each round can continue to transition from known knowledge to missing knowledge. The global model could complement its ability at each round, making its performance improve consistently.

## Relational Prototype-wise Contrastive Learning

On the server side, we calculate the global logit prototype for each class and identify their hard negative classes. Afterward, prototype-wise contrastive learning is deployed to facilitate the discrimination between classes.

At the  $t$ -th round, besides the condensed knowledge dataset, each client  $k$  uploads the logit prototypes  $\{\mathbf{p}_0^{k,t}, \dots, \mathbf{p}_{C-1}^{k,t}\}$  calculated by the global model as:

$$\mathbf{p}_c^{k,t} = \frac{1}{N_{c,k}} \sum_i^{N_{c,k}} f_{M_{t-1}}(x_{i,c,k}), \quad (11)$$

where the  $x_{i,c,k}$  denotes the local data with class  $c$  in client  $k$ , and  $f_{M_{t-1}}$  denotes the current model without the last softmax layer at the beginning of the  $t$ -th round. Then we aggregate these prototypes uploaded from each client into global logit prototypes  $\{\mathbf{p}_0^t, \dots, \mathbf{p}_{C-1}^t\}$ :

$$\mathbf{p}_c^t = \frac{1}{|\mathcal{T}_c|} \sum_k^N |\mathcal{T}_{c,k}| \mathbf{p}_c^{k,t}, \quad (12)$$

where  $N$  denotes the number of clients,  $|\mathcal{T}_c|$  denotes the total number of data of class  $c$ , and  $|\mathcal{T}_{c,k}|$  denotes the number of data of class  $c$  in client  $k$ . With global logit prototypes, we can derive the Top-K hard negative classes for class  $c$  as :

$$HN(c) = \{j_1, j_2, \dots, j_K\} = \arg \operatorname{topK}_{j \neq c} \mathbf{p}_c[j], \quad (13)$$

where  $HN(c)$  contains the class indices with Top-K values in prototype vector  $\mathbf{p}_c$  except  $c$ . We recognize  $HN(c)$  as hard negative classes' indices set for class  $c$  since the global always predicts a higher probability on these classes and tends to mis-classify. To amplify discrimination ability of the global model, it would be more effective to push features of class  $c$  away from that of  $HN(c)$ . Note that the global logit prototypes may change across rounds with the updated global model,  $HN(c)$  would also change adaptively.

We also calculate feature prototypes  $\{\mathbf{f}_0^t, \dots, \mathbf{f}_{C-1}^t\}$  with condensed knowledge datasets on the server at  $t$ -th round:

$$\mathbf{f}_c^t = \frac{1}{|\mathcal{S}_c^{t-1}|} \sum_{s_i \in \mathcal{S}_c^{0, \dots, t-1}} \psi_\theta^{t-1}(s_i), \quad (14)$$

where  $\mathcal{S}_c^{t-1}$  is the accumulated knowledge dataset of class  $c$  uploaded before  $t$ -th round, and  $\psi_\theta^{t-1}$  is the encoder of the global model  $\mathbf{M}_{t-1}$  at the very beginning of  $t$ -th round.

With hard negative classes set  $HN(c)$  and feature prototypes, inspired by SimSiam (Chen and He 2021), we propose relational supervised contrastive learning with prototypes in a bootstrap manner:

$$L_{rc} = \sum_{(s_i, c_i) \in \mathcal{S}^t} -\log \frac{\exp(h(\psi_\theta^t(s_i)) \cdot \mathbf{f}_{c_i}^t / \tau)}{\sum_{c_j \in HN(c_i)} \exp(h(\psi_\theta^t(s_i)) \cdot \mathbf{f}_{c_j}^t / \tau)}, \quad (15)$$

where  $h$  is a learnable projector similar with (Chen and He 2021; Grill et al. 2020) and  $\tau$  is the temperature hyper-parameter. Then we update the global model along with cross-entropy loss at  $t$ -th round with:

$$L_{update} = L_{ce}(\mathcal{S}^t) + L_{rc}, \quad (16)$$

where  $L_{ce}$  denotes the cross-entropy loss with condensed knowledge datasets and the relational supervised contrastive learning offers more supervision signals in model updating.

## Experiments

**Datasets.** We evaluate the performance of our proposed FedVCK on 4 medical tasks, which contain 5 datasets with different modalities from (Yang, Shi, and Ni 2021; Yang et al. 2023): 1) Colon Pathology, we adopt the Path dataset, 2) Retinal OCT scans, we adopt the OCT dataset, 3) Abdominal CT scans, we adopt the OrganS and OrganC dataset, 4) Chest X-Ray, we adopt the Pneumonia dataset. To validate the generality, we also select CIFAR10 (Krizhevsky 2009), STL10 (Coates, Ng, and Lee 2011), and ImageNette (Howard and Team 2019) datasets. Our selected datasets enjoy a wide range of modalities and resolutions from  $28 \times 28$  to  $224 \times 224$  and detailed introductions about datasets are shown in the Appendix.

**Baselines.** We compare FedVCK with nine federated learning methods including both model-centric methods (FedAvg, FedProx, and MOON) and data-centric methods (FedGen, FedMix, FedGAN, DFRD, FedDM, and DESA). We summarize rationale of the baseline selection and their synthesis objectives and methods in Table 1 in Appendix.

$\beta$	0.05					0.02				
Model	ConvNet			ResNet18		ConvNet			ResNet18	
Acc(%)	Path	OCT	OrganS	OrganC	Pneumonia	Path	OCT	OrganS	OrganC	Pneumonia
FedAvg	46.34	62.00	66.27	70.38	69.87	43.72	26.54	60.70	65.87	63.14
FedProx	61.34	62.50	69.14	71.80	69.07	40.15	32.20	67.76	60.94	62.50
MOON	51.91	56.10	52.33	71.82	62.50	50.88	29.90	62.84	61.81	62.50
FedGen	42.03	53.25	59.90	49.65	60.37	39.87	34.74	47.06	37.63	58.43
FedGAN	54.40	56.80	71.76	OOM	75.32	54.37	25.20	70.34	OOM	62.50
FedMix	35.78	48.90	62.10	60.63	62.50	31.50	29.30	54.65	29.22	62.50
DFRD	37.44	31.50	39.80	OOM	OOM	14.01	34.20	37.93	OOM	OOM
FedDM	73.97	61.70	71.37	35.60	75.80	73.64	62.20	69.46	18.20	68.75
DESA	33.37	47.00	69.98	54.16	61.38	66.41	35.20	67.32	39.46	62.50
<b>FedVCK</b>	<b>80.36</b>	<b>68.30</b>	<b>73.23</b>	<b>79.52</b>	<b>86.70</b>	<b>81.10</b>	<b>68.20</b>	<b>72.90</b>	<b>79.04</b>	<b>84.62</b>

Table 2: Overall predictive accuracy comparison on medical datasets. We test our method and baselines under two non-IID scenarios:  $Dir(0.05)$  and  $Dir(0.02)$ . For datasets with  $224 \times 224$  image sizes, we adopt the ResNet18 model. **Bold** numbers indicate the best accuracy results. ‘OOM’ indicates out-of-memory.

$\beta$	0.05					0.02				
Model	ConvNet			ResNet18		ConvNet			ResNet18	
Acc(%)	Path	OCT	OrganS	OrganC	Pneumonia	Path	OCT	OrganS	OrganC	Pneumonia
FedAvg	18.55	53.70	36.92	29.31	62.50	12.84	25.00	39.32	25.13	35.74
FedProx	56.94	49.70	48.76	50.60	63.78	40.15	28.70	49.60	46.13	62.50
MOON	49.40	32.00	45.69	33.43	62.50	26.96	25.00	42.60	25.83	62.50
FedGen	34.47	27.60	45.75	24.89	58.00	37.20	25.00	37.79	18.77	57.50
FedGAN	19.89	48.00	55.65	OOM	64.74	32.32	25.00	44.67	OOM	62.50
FedMix	28.97	31.60	57.35	32.46	62.50	28.48	25.00	42.37	17.73	62.50
DFRD	22.40	25.00	39.05	OOM	OOM	10.45	25.00	29.93	OOM	OOM
FedDM	72.76	61.70	71.32	31.66	73.27	70.39	62.20	69.33	15.12	62.50
DESA	29.48	36.90	66.00	50.72	38.78	43.62	27.80	66.42	39.46	37.50
<b>FedVCK</b>	<b>78.52</b>	<b>66.41</b>	<b>72.65</b>	<b>71.39</b>	<b>83.01</b>	<b>78.61</b>	<b>65.90</b>	<b>71.68</b>	<b>71.89</b>	<b>82.48</b>

Table 3: Predictive accuracy comparison on medical datasets under limited communication budgets.

**Configuration.** Following the commonly used setting, we simulate non-IID scenarios with Dirichlet distribution  $Dir(\beta)$  among 10 clients where  $\beta$  is 0.05 and 0.02 to simulate severe non-IID scenarios. We adopt the ConvNet (Gidaris and Komodakis 2018) and ResNet18 (He et al. 2016). We set the size of the condensed knowledge dataset  $\mathcal{S}$  to  $p\%$  of the original dataset size, where  $p$  is selected from  $\{1, 2, 5\}$  according to different datasets. We initialize  $S$  from  $\mathcal{N}(0, 1)$ . Hyper-parameters in each method are tuned as suggested in the original papers. We run all experiments with NVIDIA Geforce RTX 3090 GPU and report the mean results among three runs. More details are introduced in the Appendix.

### Performance Under Non-IID Scenarios

**Overall performance.** We evaluate all methods’ overall performance under non-IID scenarios in Table 2, assuming the communication budgets are adequate (100 communication rounds). We can observe that model-centric federated learning methods struggle with mediocre performance. Some data-centric methods (e.g. FedGen and DFRD) perform worse. We find this is because the poor synthesis quality and poor global model hinder each other and cause a vicious circle. On OrganC and Pneumonia datasets, FedGAN and DFRD face the out-of-memory problem. Clients in FedGAN must train and upload huge generators and discriminators and the server in DFRD must maintain ensemble

models and huge generators. Most data-centric baselines degrade hardly on the two datasets since capturing subtle and meaningful knowledge in larger sizes is harder. Our method successfully condenses knowledge with high quality and high necessity for the global model, thus showing advantages over all datasets’ baselines.

**Under limited communication budgets.** Since communication budgets are usually limited in reality, a method that can achieve high performance within a few communication rounds is more desired. We compare the performance of our method and baselines within 10 communication rounds under non-IID scenarios. The experimental results are shown in Table 3. We can observe that all baselines cannot achieve satisfactory performance within limited rounds, while our method consistently outperforms others on all datasets. The performance is relatively close to the overall performance in Table 2 and would not be significantly affected by more severe non-IID ( $\beta = 0.02$ ), demonstrating that our method is not only communication-efficient but also robust to non-IID.

### Performance Analysis

**Ablation study.** We conduct ablation study to evaluate the effectiveness of our designs. The experimental results are shown in Table 4. Besides, we measure the empirical MMD of the condensed knowledge between two adjacent rounds

Acc(%)	Path	OrganS	Path	OrganS
w.o. all	72.76	71.32	73.97	71.37
w.o. $L_{rc} + P_w$	73.04	71.74	76.48	72.13
w.o. $L_{rc}$	74.52	71.93	78.56	72.40
<b>FedVCK</b>	78.52	72.65	80.36	73.23

Table 4: Ablation study on medical datasets. The left part of the table is the performance under limited communication rounds and the right part is the overall performance.

Method	Path	OrganS	Pneumonia
FedMix, DRFD	12.13 MB	12.20 MB	426.15 MB
FedAvg, FedProx			
MOON, FedGen			
DESA			
FedGAN	178.85 MB	178.69 MB	2349.40 MB
FedVCK, FedDM	<b>2.04 MB</b>	<b>0.52 MB</b>	<b>11.77 MB</b>
$p\%$	1%	5%	5%

Table 5: Per-round upload communication costs.  $p\%$  indicates the size of the condensed knowledge dataset as a proportion of the size of the original dataset.

with or without model-guided selection in Figure 1b. We can note that with model-guided selection, the condensed knowledge between adjacent rounds exhibits greater MMD values, reflecting that it can avoid repeated knowledge and force the optimization process to condense more heterogeneous and model-specific knowledge. We also study the impact of the size of learnable knowledge dataset in the Appendix. Larger size would have more capacity but increase optimization difficulty and communication overhead.

**Communication analysis.** Our method is communication efficient from two aspects. From the perspective of communication rounds, we have demonstrated our method can quickly achieve satisfactory performance under limited budgets in Table 3. From the perspective of upload communication costs, we quantify the actual per-round upload communication costs of all clients in Table 5. Our method’s per-round uploading costs are less than that of model-centric federated learning. More analysis about the communication and full experimental results are shown in the Appendix.

### Privacy Analysis

To practically test whether the condensed knowledge would leak individual privacy, we conduct the membership inference attack following an advanced method: LiRA (Carlini et al. 2022) and compare FedVCK with the model-centric federated learning method (e.g. FedAvg). We attack uploaded models or condensed knowledge from clients and record the AUC of ROC on each client with a balanced test set. Since the MIA task is a binary classification, we set the minimum AUC to 0.5 and mark it as a total defense if the AUC of a client is less than or equal to 0.5. We calculate the max and mean AUC of all clients and defense rate (the proportion of clients achieving total defense) in Table 6. The results show that our method better preserves privacy than FedAvg, and enables more total defense cases. In addition,

Method	Max AUC ↓	Mean AUC ↓	Defense Rate ↑
FedAvg	0.556	0.529	10%
FedVCK	<b>0.544</b>	<b>0.514</b>	<b>50%</b>

Table 6: The AUC results of MIA experiment on OrganS dataset. ↓ means the lower, the better. ↑ means the opposite.

Acc(%)	CIFAR10		STL10		ImageNette	
$\beta$	0.05	0.02	0.05	0.02	0.05	0.02
FedAvg	52.13	52.01	46.03	39.60	47.21	38.17
FedProx	57.60	53.39	42.93	42.88	52.66	32.84
MOON	46.63	42.03	38.08	38.42	35.26	21.32
FedGen	39.36	32.71	38.11	37.44	51.92	41.89
FedGAN	55.79	53.86	51.84	50.03	50.80	40.31
FedMix	42.28	43.97	46.56	42.88	50.80	39.41
DFRD	52.07	37.53	31.60	21.03	33.20	16.20
FedDM	54.75	50.47	54.90	51.62	52.25	43.82
DESA	53.90	48.19	46.74	37.33	42.29	28.90
<b>FedVCK</b>	<b>62.96</b>	<b>60.56</b>	<b>57.04</b>	<b>56.89</b>	<b>62.76</b>	<b>61.73</b>

Table 7: Overall predictive accuracy comparison on natural datasets. We adopt the ConvNet model by default.

since our method needs fewer communication rounds, privacy can be further protected from potential temporal-based MIA (Zhu et al. 2024).

### Extend to Natural Datasets

To validate the generality of our method, we also extend our evaluation on natural datasets. The natural datasets contain various colored objects with more significant inter-class differences. The overall experimental results are shown in Table 7. Our method still outperforms others consistently, which demonstrate a boarder generality of our method. Full experiments about the predictive performance, communication cost, and privacy-preserving are listed in the Appendix.

### Conclusion and Discussion

In this paper, we propose a novel data-centric federated learning method, FedVCK, for collaborative medical image analysis. FedVCK can tackle the non-IID problem in a communication-efficient manner. Specifically, FedVCK adaptively selects the most necessary knowledge with the guidance of current models, and condenses it into a small knowledge dataset with latent distribution constraints to enhance the quality. The condensed knowledge can effectively update the global model with the help of relational supervised contrastive learning. Our method generally outperforms state-of-the-art methods in non-IID scenarios, especially under limited communication budgets. Further work is to extend to more data modalities such as 3D CT and to adopt advanced techniques to improve the effectiveness and efficiency of condensation.

### Acknowledgments

This work is supported by the National Key R&D Program of China under Grant No.2022YFB2703301.

## 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*.
- Carlini, N.; Chien, S.; Nasr, M.; Song, S.; Terzis, A.; and Tramer, F. 2022. Membership inference attacks from first principles. In *2022 IEEE Symposium on Security and Privacy (SP)*, 1897–1914. IEEE.
- Chen, H.-Y.; and Chao, W.-L. 2020. Fedbe: Making bayesian model ensemble applicable to federated learning. *arXiv preprint arXiv:2009.01974*.
- Chen, X.; and He, K. 2021. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15750–15758.
- Coates, A.; Ng, A.; and Lee, H. 2011. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, 215–223. JMLR Workshop and Conference Proceedings.
- Dong, T.; Zhao, B.; and Lyu, L. 2022. Privacy for free: How does dataset condensation help privacy? In *International Conference on Machine Learning*, 5378–5396. PMLR.
- Gao, X.; Chen, T.; Zhang, W.; Li, Y.; Sun, X.; and Yin, H. 2024a. Graph condensation for open-world graph learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 851–862.
- Gao, X.; Yu, J.; Chen, T.; Ye, G.; Zhang, W.; and Yin, H. 2024b. Graph condensation: A survey. *arXiv preprint arXiv:2401.11720*.
- Gidaris, S.; and Komodakis, N. 2018. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4367–4375.
- Gretton, A.; Borgwardt, K.; Rasch, M.; Schölkopf, B.; and Smola, A. 2012. A kernel two-sample test. *Journal of Machine Learning Research, Journal of Machine Learning Research*.
- Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33: 21271–21284.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Howard, J.; and Team, F. 2019. Imagenette: A subset of 10 easily classified classes from Imagenet. Available at <https://github.com/fastai/imagenette>. Accessed on 2024-08-01.
- Huang, C.-Y.; Srinivas, K.; Zhang, X.; and Li, X. 2024. Overcoming Data and Model Heterogeneities in Decentralized Federated Learning via Synthetic Anchors. *arXiv preprint arXiv:2405.11525*.
- Ioffe, S.; and Szegedy, C. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, 448–456. pmlr.
- Krizhevsky, A. 2009. Learning multiple layers of features from tiny images. Technical Report, University of Toronto. CIFAR-10 dataset available at <https://www.cs.toronto.edu/~kriz/cifar.html>.
- Li, Q.; Diao, Y.; Chen, Q.; and He, B. 2022a. Federated learning on non-iid data silos: An experimental study. In *2022 IEEE 38th international conference on data engineering (ICDE)*, 965–978. IEEE.
- Li, Q.; He, B.; and Song, D. 2021. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10713–10722.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.
- Li, X.; Huang, K.; Yang, W.; Wang, S.; and Zhang, Z. 2019. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*.
- Li, Z.; Shao, J.; Mao, Y.; Wang, J. H.; and Zhang, J. 2022b. Federated learning with gan-based data synthesis for non-iid clients. In *International Workshop on Trustworthy Federated Learning*, 17–32. Springer.
- Lin, T.; Kong, L.; Stich, S. U.; and Jaggi, M. 2020. Ensemble distillation for robust model fusion in federated learning. *Advances in neural information processing systems*, 33: 2351–2363.
- Mothukuri, V.; Parizi, R. M.; Pouriyeh, S.; Huang, Y.; Dehghantaha, A.; and Srivastava, G. 2021. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115: 619–640.
- Nguyen, D. C.; Ding, M.; Pathirana, P. N.; Seneviratne, A.; and Zomaya, A. Y. 2021. Federated learning for COVID-19 detection with generative adversarial networks in edge cloud computing. *IEEE Internet of Things Journal*, 9(12): 10257–10271.
- Verma, V.; Lamb, A.; Beckham, C.; Najafi, A.; Mitliagkas, I.; Lopez-Paz, D.; and Bengio, Y. 2019. Manifold mixup: Better representations by interpolating hidden states. In *International conference on machine learning*, 6438–6447. PMLR.
- Wang, S.; Fu, Y.; Li, X.; Lan, Y.; Gao, M.; et al. 2024a. DFRD: Data-Free Robustness Distillation for Heterogeneous Federated Learning. *Advances in Neural Information Processing Systems*, 36.
- Wang, T.; Zhu, J.-Y.; Torralba, A.; and Efros, A. A. 2018. Dataset distillation. *arXiv preprint arXiv:1811.10959*.
- Wang, Y.; Fu, H.; Kanagavelu, R.; Wei, Q.; Liu, Y.; and Goh, R. S. M. 2024b. An aggregation-free federated learning for tackling data heterogeneity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26233–26242.

- Xie, L.; Lin, M.; Liu, S.; Xu, C.; Luan, T.; Li, C.; Fang, Y.; Shen, Q.; and Wu, Z. 2024a. pFLFE: Cross-silo Personalized Federated Learning via Feature Enhancement on Medical Image Segmentation. In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15010. Springer Nature Switzerland.
- Xie, L.; Lin, M.; Luan, T.; Li, C.; Fang, Y.; Shen, Q.; and Wu, Z. 2024b. MH-pFLID: Model Heterogeneous personalized Federated Learning via Injection and Distillation for Medical Data Analysis. In *Forty-first International Conference on Machine Learning*.
- Xie, L.; Lin, M.; Xu, C.; Luan, T.; Zeng, Z.; Qian, W.; Li, C.; Fang, Y.; Shen, Q.; and Wu, Z. 2024c. MH-pFLGB: Model Heterogeneous personalized Federated Learning via Global Bypass for Medical Image Analysis. In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15010. Springer Nature Switzerland.
- Xiong, Y.; Wang, R.; Cheng, M.; Yu, F.; and Hsieh, C.-J. 2023. Feddm: Iterative distribution matching for communication-efficient federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16323–16332.
- Yang, J.; Shi, R.; and Ni, B. 2021. MedMNIST Classification Decathlon: A Lightweight AutoML Benchmark for Medical Image Analysis. In *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 191–195.
- Yang, J.; Shi, R.; Wei, D.; Liu, Z.; Zhao, L.; Ke, B.; Pfister, H.; and Ni, B. 2023. MedMNIST v2-A large-scale lightweight benchmark for 2D and 3D biomedical image classification. *Scientific Data*, 10(1): 41.
- Yin, H.; Molchanov, P.; Alvarez, J. M.; Li, Z.; Mallya, A.; Hoiem, D.; Jha, N. K.; and Kautz, J. 2020. Dreaming to distill: Data-free knowledge transfer via deepinversion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8715–8724.
- Yoon, T.; Shin, S.; Hwang, S. J.; and Yang, E. 2021. Fedmix: Approximation of mixup under mean augmented federated learning. *arXiv preprint arXiv:2107.00233*.
- Yu, R.; Liu, S.; and Wang, X. 2023. Dataset distillation: A comprehensive review. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, H.; Li, S.; Wang, P.; Zeng, D.; and Ge, S. 2024. M3D: Dataset Condensation by Minimizing Maximum Mean Discrepancy. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, 9314–9322. AAAI Press.
- Zhang, L.; Shen, L.; Ding, L.; Tao, D.; and Duan, L.-Y. 2022. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10174–10183.
- Zhao, B.; and Bilen, H. 2023. Dataset Condensation with Distribution Matching. In *2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*.
- Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; and Chandra, V. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.
- Zheng, X.; Ying, S.; Zheng, F.; Yin, J.; Zheng, L.; Chen, C.; and Dong, F. 2023. Federated Learning on Non-iid Data via Local and Global Distillation. In *2023 IEEE International Conference on Web Services (ICWS)*, 647–657. IEEE.
- Zhou, T.; Zhang, J.; and Tsang, D. H. 2023. FedFA: Federated learning with feature anchors to align features and classifiers for heterogeneous data. *IEEE Transactions on Mobile Computing*.
- Zhu, G.; Li, D.; Gu, H.; Han, Y.; Yao, Y.; Fan, L.; and Yang, Q. 2024. Evaluating Membership Inference Attacks and Defenses in Federated Learning. *arXiv preprint arXiv:2402.06289*.
- Zhu, L.; Liu, Z.; and Han, S. 2019. Deep leakage from gradients. *Advances in neural information processing systems*, 32.
- Zhu, W.; and Luo, J. 2022. Federated medical image analysis with virtual sample synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 728–738. Springer.
- Zhu, Z.; Hong, J.; and Zhou, J. 2021. Data-free knowledge distillation for heterogeneous federated learning. In *International conference on machine learning*, 12878–12889. PMLR.