

CD²-pFed: Cyclic Distillation-guided Channel Decoupling for Model Personalization in Federated Learning

Yiqing Shen¹, Yuyin Zhou², Lequan Yu^{3*}

¹ Shanghai Jiao Tong University, ² UC Santa Cruz, ³ The University of Hong Kong

shenyq@sjtu.edu.cn, zhouyuyiner@gmail.com, lqyu@hku.hk

Abstract

Federated learning (FL) is a distributed learning paradigm that enables multiple clients to collaboratively learn a shared global model. Despite the recent progress, it remains challenging to deal with heterogeneous data clients, as the discrepant data distributions usually prevent the global model from delivering good generalization ability on each participating client. In this paper, we propose **CD²-pFed**, a novel **Cyclic Distillation-guided Channel Decoupling** framework, to personalize the global model in FL, under various settings of data heterogeneity. Different from previous works which establish layer-wise personalization to overcome the non-IID data across different clients, we make the first attempt at channel-wise assignment for model personalization, referred to as channel decoupling. To further facilitate the collaboration between private and shared weights, we propose a novel cyclic distillation scheme to impose a consistent regularization between the local and global model representations during the federation. Guided by the cyclical distillation, our channel decoupling framework can deliver more accurate and generalized results for different kinds of heterogeneity, such as feature skew, label distribution skew, and concept shift. Comprehensive experiments on four benchmarks, including natural image and medical image analysis tasks, demonstrate the consistent effectiveness of our method on both local and external validations.

1. Introduction

Deep learning techniques have received notable attention in various vision tasks, such as image classification [7], object detection [31], and semantic segmentation [25]. Yet, the success of deep neural networks heavily relies on a tremendous volume of valuable training images. One possible solution is to collaboratively curate numerous data samples from different parties (e.g., different mobile devices

*Corresponding Author.

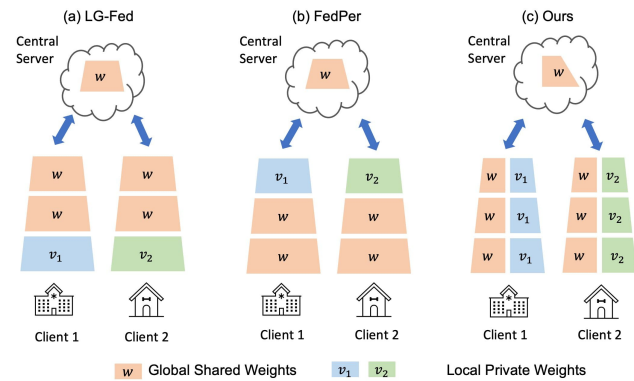


Figure 1. Illustration of different parameter decoupling manners for model personalization in Federated Learning. The previous approaches combine local and global parameters in a layer-wise mechanism, including LG-Fed [22] in low-level input layers (a) and FedPer [2] in high-level output layers (b). Instead, we achieve model personalization via channel-wise decoupling (c).

and companies). However, collecting distributed data into a centralized storage facility is costly and time-consuming. Additionally, in real practice, decentralized image data should not be directly shared, due to privacy concerns or legal restrictions [1, 39]. In this case, conventional centralized machine learning frameworks fail to satisfy the data privacy protection constraint.

Therefore, the data-private distributed training paradigms, especially Federated Learning (FL), have received an increasing popularity [3–5, 19, 24, 28, 36, 47, 50]. To be more specific, in FL, a shared model is globally trained with an orchestration of local updates within data stored at each client. A pioneering FL algorithm named Federated Average (FedAVG), aggregates parameters at the central server by communication across clients once per global epoch, without explicit data sharing [28]. Compared with local training, the federation on a larger scale of training data has demonstrated its superiority to boost the generalization ability on unseen data, with the orchestration of distributed private data [6, 28].

However, data heterogeneity is one of the most fundamental challenges faced by FL. The concept of independent and identically distributed (IID) is clear, while data can be non-IID in many ways, *e.g.*, feature skew, label distribution skew, or concept shift [11]. Previously, sharp performance degradation was observed on FedAVG with unbalanced and non-IID data. This ill-effect is attributed to the weight divergence, which can be quantified by the earth mover’s distance between distributions over classes [49]. Although each client can train a private model locally by optimizing the objective with no information change among each other, it would inevitably result in overfitting and a poor generalization ability on new samples. As suggested in [28], simply sharing a small subset of data globally greatly enhances the generalization of FedAVG. However, this scheme cannot be directly applied to real-world tasks due to the violation of privacy concerns.

Consequently, researchers have sought to train a collection of models that is stylized for each local distribution to enable stronger performance for each participating client without requiring any data sharing [49], which is known as personalized federated learning PFL [37]. Various approaches have been proposed to accomplish the model personalization in FL [6, 10, 35, 37]. Among these different paradigms, one popular solution is to directly assign personalized parameters for each local client. For this line of methods, the private personalized parameters are trained locally and not shared with the central server. Existing works have made attempts to achieve personalization by assigning personalized parameters in either top layers [2] or bottom layers [22]. However, these approaches usually require prior knowledge for the determination of which layers to be personalized. More critically, we observe performance degradation that existing PFL approaches fail to achieve a consistent generalization over comprehensive settings of data heterogeneity [29]. Additionally, existing layer-wise personalization approaches cannot effectively handle the discrepancy between the learned local and global model representations due to the weight divergence [49]. The inferior performance of some local clients motivates us to seek a more generic yet efficient combination between the local and global information.

In light of these challenges, we propose **CD²-pFed**, a novel Cyclic Distillation-guided Channel Decoupling framework for model personalization in FL. As shown in Figure 1, different from previous layer-wise personalization approaches, *e.g.*, FedPer [2] and LG-Fed [22], the proposed novel *channel decoupling* paradigm dynamically decouples the parameters at the channel dimension for personalization instead. By employing learnable personalized weights at all layers, our channel decoupling paradigm no longer requires heuristics for designing specific personalization layers. More importantly, our method achieves model per-

sonalization for both low-level and high-level layers, which facilitates tackling feature heterogeneity, distribution skew, and concept shift.

To bridge the semantic gap between the learned visual representation from the decoupled channels, we further propose a novel *cyclic distillation* scheme by mutually distilling the local and global model representation (*i.e.*, soft predictions by the private and shared weights) from each other. Benefiting from the distilled knowledge, our channel decoupling framework enables synergistic information exchange between the global and local model training, therefore preventing biased local model training on non-IID data. Extensive experimental results on both heterogeneous data and exterior unseen samples [22] demonstrate that our method largely improves the generalization of FedAVG with negligible additional computation overhead. Below, we summarize the major contributions of this work.

- We propose a novel *channel decoupling* paradigm to decouple the global model at the channel dimension for personalization. Instead of using personalization layers for tackling either feature or label distribution skew, our approach provides a unified solution to address a broad range of data heterogeneity.
- To further enhance the collaboration between private and shared weights in channel decoupling, we design a novel *cyclic distillation* scheme to narrow the divergence between them.
- We compare our method with previous state-of-the-art PFL approaches on four benchmark datasets, including synthesized and real-world image classification tasks, with different kinds of heterogeneity. Results demonstrate the superiority of our method over state-of-the-art PFL approaches.

2. Related Work

2.1. Data Heterogeneity

Federated Averaging (FedAVG) is a prevailing federated learning algorithm to train a global model with distributed data [28]. Under the assumption of unbalanced and IID (independent and identically distributed) characteristics, FedAvg has achieved notable empirical success on its robustness and performance. However, facing the variety and diversity of real-world data, the condition of non-IID is unrealistic to be ensured. Instead, statistical data heterogeneity is a more general case, where the data distribution of local clients deviates significantly from the global distribution. It results in sharp performance degradation, *e.g.*, up to 11.0% for MNIST and 51.0% for CIFAR-10 [49], where the common predictor does not generalize well on local data. One of the most fundamental challenges in training a robust FL

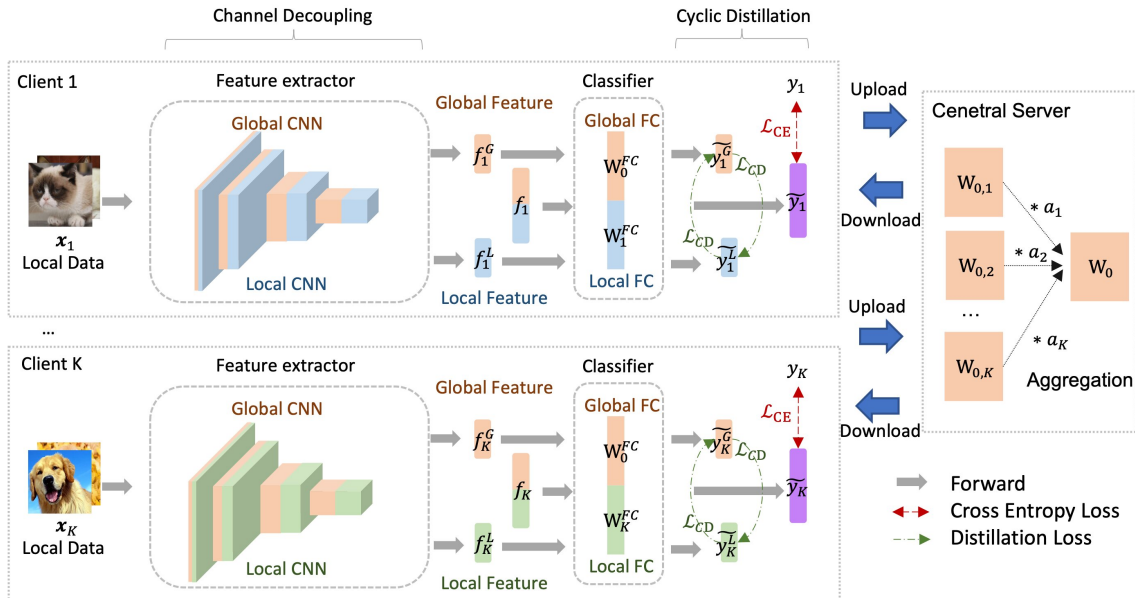


Figure 2. A schematic illustration of our proposed CD^2 -pFed framework for model personalization in federated learning. We use the blue and green to mark out personalized channels and features, which reside locally; and the orange for global representations.

is the presence of non-IID data. Concretely, the underlying distribution for arbitrary paired clients is very likely to differ from each other. There are various formulations on the existence of non-IID, including the feature distribution skew, label distribution skew, and the concept shift [11]. To tackle the data heterogeneity, Li *et al.* proposed an optimization scheme, namely FedProx, to re-parametrize FedAvg with variable amounts of work to be performed locally across devices [20]. FedBN employs local batch normalization to alleviate the feature distribution skew before the model aggregation [21].

2.2. Personalized Federated Learning

Three major challenges that restrict the generalization ability of federated global model on local data are 1) device heterogeneity, 2) data heterogeneity due to the non-IID distribution, 3) model heterogeneity to adapt to local environment [40]. Among those, data-IID is the most practical issue. Yet, the privacy protection mechanisms in conventional FL conflict with achieving higher performance for each individual user [44]. Subsequently, Personalized Federated Learning (PFL) has received numerous attentions from researchers to cope with the above three challenges. In PFL, the global model is personalized for each local client and plays an intermediate paradigm between pure local training and FL [27]. Leveraging a personalized model per client, PFL can integrate the client’s own dataset and orchestration of data from other clients into the training process.

There are various techniques to adapt the global model for personalization [15], including transfer learning [33],

multi-task learning [35], meta-learning [10], knowledge distillation [45], and network decoupling [2,22]. We mainly focus on the network decoupling methods, where the global network is decoupled into personalized layers which reside locally, and global layers. For example, FedPer splits a neural network architecture into base layers, which are trained centrally by FedAvg, and the top personalized layers [2], which are trained separately with the personalized layers. Very similarly to FedPer, LG-Fed personalizes the bottom layer, while keeping the top layers shared across all involved clients [22]. FedPer shows its superiority with an observable labeling skew such as on FLICKR-AES [30], while LG-Fed with data skews such as CIFAR non-IID split. However, both skews exist in real-world tasks. Therefore, in this work, we attempt to bridge the gap between top layers personalization and bottom layers personalization with a unified channel decoupling scheme.

2.3. Knowledge Distillation

The key idea of knowledge distillation (KD) is to transfer the dark knowledge from a pre-trained teacher model to a lightweight student network by learning its soft predictions, intermediate features or attention maps [9,29,42]. KD has broad applications in machine learning and computer visions fields, including transfer learning, semi-supervised learning, reinforcement learning [13,32,38,41,48]. KD has also achieved remarkable performance as a regularization scheme. For example, Yun *et al.* [46] distilled predictive distribution between samples of the same label to mitigate overconfident predictions and reduce intra-class variations

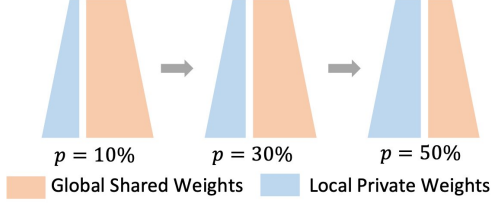


Figure 3. Illustration of Channel Decoupling with progressive model personalization.

in an image classification task. In FL scenarios, an ensemble distillation scheme was proposed to replace the aggregation for model fusion [23]. Besides, Li *et al.* [18] adopted knowledge distillation to personalize the global model by translating knowledge between participants.

3. Methodology

3.1. Problem Formulation

We consider a set of K clients, which are all connected to a central server. Moreover, each client only has access to its local data, denoted as \mathcal{D}_i , with no data sharing between clients. In a data heterogeneous setting, the underlying distribution of \mathcal{D}_i , denoted as \mathcal{P}_i are not identical, *i.e.*, $\mathcal{P}_i \neq \mathcal{P}_j$. Specifically, there are three common categories to depict the non-IID characteristic in FL [11]: 1) feature distribution skew (covariate shift), *i.e.*, $\mathcal{P}_i(x) \neq \mathcal{P}_j(x)$; 2) label distribution skew (prior probability shift), *i.e.*, $\mathcal{P}_i(y) \neq \mathcal{P}_j(y)$; and 3) same label but different features (concept drift), *i.e.*, $\mathcal{P}_i(x|y) \neq \mathcal{P}_j(x|y)$.

The goal of our work is to train a collection of K models to adapt to the local dataset without exchanging their local data with other parties. The network at the i -th client ($i \in \{1, \dots, K\}$) is composed of private personalized parameters w_i , and global shared parameters w_0 . More formally, we formulate the loss function corresponding to the i -th client as $F_i: \mathbb{R}^d \rightarrow \mathbb{R}$, and then the overall objective in personalized federated learning is defined as follows,

$$\min_{\{w_i\}_{i=0}^K} F(w_0, w_1, \dots, w_K) = \sum_{i=1}^K \alpha_i \cdot F_i(w_0, w_K). \quad (1)$$

The balancing weight α_i depends on the scale of the private dataset, *i.e.*, $\alpha_i = \frac{|\mathcal{D}_i|}{\sum_j |\mathcal{D}_j|}$. In this scenario, we consider supervised learning, leading to

$$F_i(w_0, w_i) = \mathbb{E}_{(\mathbf{x}_j, y_j) \sim \mathcal{P}_i} [l_i(\mathbf{x}_j, y_j; w_0, w_i)], \quad (2)$$

where l_i measures the sample-wise loss between the prediction of the network parameterized by (w_0, w_i) and the ground truth label y_j when given the input image \mathbf{x}_i .

3.2. Channel Decoupling for Model Personalization

As shown in Figure 2, we propose a vertically channel-wise decoupling framework to personalize the global model for non-IID federated learning. Concretely, we assign an adaptive proportion of learnable personalized weights at each layer from the target model, from top to bottom. In this fashion, our framework is intended to achieve higher personalization capacity for both simple and complex patterns (*e.g.*, image and label-level personalization). We define a uniform personalization partition rate $p \in [0, 1]$ to determine the precise proportion of the personalized channels in each layer. It follows that these p proportion of channel parameters are trained locally, without the aggregation by the central server. Subsequently, these private weights vary from each other, written as w_i where the subscript is associated with its client ID. The remaining $(1 - p)$ percentage of the shared weights, denoted as w_0 , are trained with common FL algorithms, such as FedAvg [28].

A larger value in p represents a higher degree of personalization. Therefore, the case when $p = 0$ degenerates into the conventional FedAvg [28] with no model personalization, and conversely $p = 1$ denotes a full local training procedure in an absence of the federated communications. One significant benefit from our vertical decoupling strategy, compared with horizontally layer-wise personalization ones, such as LG-Fed [22] and Fed-Per [2], is to enable a model personalization from lower to topper layers, resulting in a potential to a more general framework for weights personalization, as well as improving its capacity to address a broader range of data heterogeneity, such as both feature and label distribution heterogeneity.

Progressive model personalization. One key element in our proposed channel decoupling scheme is the determination of the personalized ratio p in each layer. As aforementioned, the personalized ratio p controls the volume of private weights to learn the local representation, which determines the capacity to learn a good representation on the heterogeneous data. We are motivated to provide a better initialization for the personalization from the globally learned representation. Thus, as shown in Figure 3, instead of a fixed p , the model capability to learn local personalized features is taken into account by a progressive increment scheme. That is, in the initial stage, we set p to a small value to facilitate learning global representation for faster convergence, and afterwards gradually increase its value. Consequently, we increase the value of p progressively depending on the global epoch number T . Similar to the learning rate schedule, a variety of schemes exist for the increment, such as cyclical learning rate [34], cosine annealing [26]. For simplicity, here we apply the linear growth scheme, *i.e.*,

$$p_t = p \cdot \frac{t}{T}, \quad (3)$$

where T is the total global epoch number, t is the current epoch number, and p is the maximum personalization ratio.

3.3. Cyclic Distillation

The backbone neural network is decoupled into personalized weights w_i and shared weights w_0 respectively, and afterwards trained simultaneously by optimizing the local supervised loss in Eq. (2). However, as the local personalized and global parameters are learned with different distribution data, the statics of these parameters suffer from divergence and subsequently lead to the performance degradation [49]. Additionally, explicit consistency regularization between two parts is absent, during the optimization of local supervised objectives in most previous works. To cope with this issue, we make the first attempt at introducing self distillation into PFL to improve the client-side weights inner communications between the private and shared model weights and thus reduce the gap between the learned representations from local and global weights. Motivated by inplace distillation [43], the key idea in the proposed Cyclic Distillation is to impose a consistency regularization between w_i and w_0 in the local training procedure, as depicted in Figure 2.

We write the subnet parameterized by w_i, w_0 as g_{w_i}, g_{w_0} , and the network composed of (w_i, w_0) as g_{w_i, w_0} . Notably, g_{w_i} intends to learn personalized local representation from \mathcal{D}_i , whereas g_{w_0} to learn global general representation. For each input sample \mathbf{x}_i , we collect the predictions $\tilde{y}_i, \tilde{y}_i^L, \tilde{y}_i^G$ from $g_{w_i, w_0}, g_{w_i}, g_{w_0}$, respectively. The overall prediction \tilde{y}_i minimizes the cross entropy loss \mathcal{L}_{CE} with ground truth y_i . The cyclic distillation loss is defined as:

$$\mathcal{L}_{CD} = \frac{1}{2} (KL(\tilde{y}_i^L, \tilde{y}_i^G) + KL(\tilde{y}_i^G, \tilde{y}_i^L)), \quad (4)$$

where $KL(\cdot, \cdot)$ denotes the Kullback-Leibler (KL) divergence. It can impose a consistency regularization between w_i and w_0 , guiding the predictions from w_i and w_0 to align with each other. Consequently, the overall loss function is

$$\mathcal{L} = \mathcal{L}_{CE} + \lambda \cdot \mathcal{L}_{CD}, \quad (5)$$

with the balancing coefficient λ set to 1 in this work.

Temporal average moving for personalized weights. To stabilize the training performance, we utilize an exponential moving average (EMA) scheme for the local weights update in personalized channel w_i for a more smoothing training dynamics. We use the superscript l to mark the corresponded local epoch number, then the EMA update of w_i at t is

$$w_i^l = \beta_t w_i^l + (1 - \beta_t) w_i^{l-1}, \quad (6)$$

where w_i^l is the raw update from Eq. (5). The smoothing coefficient β_t depends on the current global epoch number

Algorithm 1 Local training with CD²-pFed at client i .

Input: local epoch number η_i

Output: w_0^t

- 1: Download w_0^{t-1} from Central Server
 - 2: Update Personalized Ratio p
 - 3: **for do** $l = 1, 2, \dots, \eta_i$
 - 4: Sample of Batch of Data from \mathcal{D}_i
 - 5: Forward and Compute Cross Entropy Loss \mathcal{L}_{CE}
 - 6: Compute Cyclic Distillation Loss \mathcal{L}_{CD} in Eq. (4)
 - 7: Update the Weights
 - 8: Adjust Personalized Weights w_i^l with Eq. (6)
 - 9: **end for**
 - 10: Upload w_0^t to Central Server
-

and follows a ramp-up strategy in previous works [16] *i.e.*,

$$\beta_t = \begin{cases} \beta \cdot \exp(-5(1 - \frac{t}{t_0})^2), & t \leq t_0 \\ \beta, & t > t_0 \end{cases}, \quad (7)$$

where β is set to 0.5, and t_0 is set to 10% of the total federated epoch numbers.

Method overview. We summarize the thorough local training procedure at each client in Alg. 1. Afterwards, the central server collects all global weights w_0 from each client and adopts FedAvg to aggregate them.

4. Experiments

4.1. Datasets

Focusing on image classification tasks, we use four benchmark datasets to evaluate the proposed CD²-pFed, namely CIFAR-10, CIFAR-100, FLICKR-AES, and a combination of public and private histology images, termed as HISTO-FED in this paper.

CIFAR-10 contains a total number of 60000 color images sized at 32×32 in 10 classes, with 5000 training images and 1000 test images per class [14]. We focus on a highly non-IID setting, *i.e.* characterized as label distribution skew. We follow previous works [22, 28] to assign images from at most $s \in \{2, 3, 4, 5, 8, 10\}$ classes to each client. A higher s corresponds to higher variance in data distribution. For example, $s = 10$ is an IID setting, while $s = 2$ is the highest heterogeneous data split. We set the client number $K = 10$ for CIFAR-10, as the literature [2, 22].

CIFAR-100 contains 500 training images and 100 testing images per class, with a total number of 100 classes [14]. Similar to CIFAR-10, the color images are scaled at 32×32 . We set the client number $K = 30$ and assign at most $s = 40$ classes to each client [2], which is also non-IID (*i.e.* label distribution skew).

FLICKR-AES is used to evaluate the performance of personalized image aesthetics in many literature [30]. The images are randomly split to 80% for training and 20% for testing. Additionally, REAL-CUR is leveraged as an external test set to evaluate the global model representation in the context of real-world personal photo ranking. Images from 14 personal albums, with an average of 197 to 222 images per album, were collected and rated by one user [30]. Due to the personal bias in aesthetic scoring, the non-IID is characterized as concept shift, leading to non-IID data distribution. We use a subset of $K = 30$ users as clients, the same as the setting in previous work [2].

HISTO-FED is a medical image datasets, consisting of both public and private hematoxylin & eosin (H&E) stained histological whole-slide images of human colorectal cancer (CRC) and normal tissue. They are curated from four medical centers. We set the client number $K = 3$ where each of them uses images from one medical center, and the remaining center is used as the external test set. Client 1 and client 2 use subsets of total slides number $N = 86$ and 50 from two public datasets NCT-CRC-HE-100K, CRC-VAL-HE-7K respectively. Each of them has a number of 7180 image patches, spitted from slides. Client 3 and external test set have 7000, and 4000 image patches respectively, curated from a private dataset of slides number $N = 20$ and 10 . Each image is labeled with one of nine categories. All images involved in this research received appropriate ethical approval. Due to the stain variance [12], the images between clients are highly non-IID, depicted as the feature skew. (as illustrated in Figure 4).

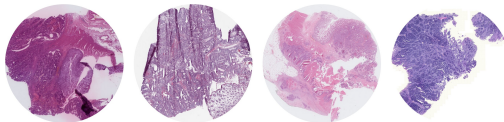


Figure 4. Illustrate examples of stain variant histology whole slides from different medical centers.

4.2. Experimental Settings

Backbone Architectures. We adopt the following network backbones for performance evaluation: 1) LeNet-5 [17] for CIFAR-10, 2) ResNet-34 [8] for CIFAR-100, FLICKR-AES, and 3) ResNet-32 for HISTO-FED, following the previous works [2, 12, 22]. All backbone networks are trained from scratch, without loading any pre-trained weights.

Hyper Parameters. At each local client, we employ stochastic gradient descent optimizer where the Nesterov momentum and the weight decay rate are set to 0.9 and 5×10^{-4} respectively. The local epoch number $\eta_i = 1$, batch size $b = 128$ for CIFAR-10; $\eta_i = 4$, $b = 128$ for CIFAR-100; $\eta_i = 4$, $b = 4$ for FLICKR-AES and HISTO-

FED. The total epoch number T is set to 50. Additionally, in CD²-pFed, we set maximum smoothing coefficient for EMA $\alpha = 0.5$ in Eq. (7), balancing coefficient in loss function $\lambda = 1$ in Eq. (5), and $p = 0.5$ on CIFAR-10/100 and HISTO-FED and $p = 0.8$ on FLICKR-AES in Eq. (3).

Comparison Methods. We compare our CD²-pFed with FedAvg [28], local training, LG-Fed [22] and FedPer [2]. FedAvg is the conventional federated learning algorithm, where no personalization is involved [28]. Local training trains a collection of K models for each client, without communications between clients. Personalized Federated Learning models play an intermediate role between FedAvg and local training, by personalizing the global structure with local unshared weights. In this work, we primarily compared our method with a model modification-based personalization scheme, which is more similar to ours. LG-Fed jointly learns compact local representations on each device with lower layers and a global model in top layers across all clients [22]. FedPer designed a base plus top personalization layer structure, conversely to LG-Fed which assigns personalization to bottom layers.

Implementations. All experiments are conducted on one NVIDIA Tesla V100 GPU with 32Gb memory. The proposed CD²-pFed is implemented on Pytorch 1.6.0 in Python 3.7.0 environment. We used an public implementation for FedAvg, LG-Fed, and FedPer for comparison. All the extra hyper-parameters involved in the compared methods are retained as their original settings.

Evaluation Metrics. We use two metrics on CIFAR-10, and CIFAR-100, following previous work [22]. 1) *Local Test Top-1 Classification Accuracy (%)*. We know precisely the client where the data sample belongs, thus we can choose the particular trained local model to predict. It evaluates the performance of model personalization. 2) *New Test Top-1 Classification Accuracy (%)*. We do not know the client where the data sample belongs to, thus we employ an ensemble of all local models to derive averaged predictions, where the local model will be uploaded to the central server. This index measures the compatibility between local and global model representation. On FLICKR-AES and HISTO-FED, we use one more index, namely the *External Test Top-1 Classification Accuracy (%)*. Specifically, we use external test samples in addition to the local or new test. Thus, these images potentially are potentially sampled from different distributions, intended to verify the generalization on the global model representation. Notably, we do not perform the external validation on CIFAR-10/100 due to the absence of external samples.

4.3. Experimental Results on Synthesized Data

Effect of Data Heterogeneity. We first evaluate the performance on CIFAR-10 on different levels of data heterogene-

ity, which is quantified by s . As shown in Figure 5, on all degree of heterogeneity, *i.e.*, s , CD²-pFed consistently outperforms LG-Fed and FedPer. Moreover, the performance gap monotonically increases with the heterogeneity. When $s = 10$, *i.e.*, in a IID setting, achieves very marginally the same test accuracy. In the rest of this section, we focus on the most non-IID settings.

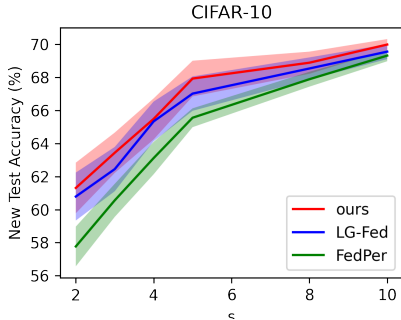


Figure 5. Effect of data heterogeneity on model personalization.

Results on CIFAR-10. As shown in Table 1, our proposed PFL frameworks significantly improve the backbone network by 31.83%, on the highest degree of data heterogeneity, *i.e.*, $s = 2$. This empirical success shows the effectiveness of personalization by the channel-wise ensemble. Additionally, compared with the state-of-the-art layer-wise personalization scheme [2, 22], our approach achieves both the best local and new classification accuracy, suggesting our model learns more capable local and global representation. Additionally, the superiority of the new test accuracy suggests that our scheme also achieves higher generalization on unseen data in personalization, attributed to the equal role of personalized and shared weights they played in FL.

Table 1. Comparison of personalized federated learning methods on CIFAR-10 with a highest heterogeneity non-IID split, *i.e.*, $s = 2$. The best results are marked in **bold**, and results reported in [22] are indicated by *. Two metrics namely the local test and new test classification accuracy, are used for evaluating the model personalization and generalization respectively.

Methods	Local (\uparrow)	New (\uparrow)
FedAvg [28]*	58.99 \pm 1.50	58.99 \pm 1.50
Local Train*	87.92 \pm 2.14	10.03 \pm 0.06
LG-Fed [22]*	91.77 \pm 0.56	60.79 \pm 1.45
FedPer [2]	83.29 \pm 0.98	57.77 \pm 1.98
Ours	91.82\pm0.43	61.31\pm1.53

Results on CIFAR-100. As illustrated in Table 2, a local test accuracy improvement of 28.75% is achieved with CD²-pFed, showing its effectiveness on model personalization for the local dataset with more wealthy categories.

Meanwhile, there is a 5.92 new test accuracy improvement, yielding its generalization on unseen data. We also outperform layer-wise personalization methods such as FedPer, LG-Fed.

Table 2. Comparison of personalized federated learning methods on CIFAR-100 with non-IID split $s = 40$. The best results are marked in **bold**.

Methods	Local (\uparrow)	New (\uparrow)
FedAvg [28]	29.23 \pm 1.75	29.23 \pm 1.75
Local Train	44.59 \pm 0.90	11.98 \pm 0.22
LG-Fed [22]	56.77 \pm 0.75	34.50 \pm 1.02
FedPer [2]	53.24 \pm 2.33	30.47 \pm 1.73
Ours	57.98\pm0.64	35.15\pm0.56

4.4. Experimental Results on Real-world Data

Results on FLICKR-AES. We test the local training performance on FLICKR-AES due to the small scale of local clients, making it easy to suffer from overfitting. In, FLICKR-AES, the labeling distribution is non-IID, fitting the philosophy of FedPer. With a marginal outperform to LG-FED, FedPer shows its superiority of top layer personalization in tackling with label distribution skew, while LG-Fed suffers an inferior performance. It is worth noticing that LG-Fed only slightly outperforms the baseline FedAvg, this is attributed to the fact that skew exists in label distributions where the personalized bottom layers in LG-Fed are difficult to learn. The empirical comparison is summarized in Table 3, where our framework outperforms state-of-the-art personalization schemes on both local and external tests. Additionally, CD²-pFed can significantly reduce the test variance, leading to more stable and robust predictions. Conclusively, CD²-pFed equipped with both top and bottom personalization does not suffer from the ill-effect on LG-Fed in facing label distribution skew.

Table 3. Comparison of personalized federated learning methods on FLICKR-AES, and the external validation of REAL-CUR. The best results are marked in **bold**.

Methods	Local (\uparrow)	External (\uparrow)
FedAvg [28]	24.50 \pm 2.01	20.08 \pm 1.34
LG-Fed [22]	25.78 \pm 2.40	20.98 \pm 1.34
FedPer [2]	43.26 \pm 3.23	40.55 \pm 1.78
Ours	47.89\pm2.03	45.67\pm1.67

Results on HISTO-FED. The internal and external test results of CD²-pFed on four clients consistently outperform the baselines, in Table 4. We achieve higher improvement from the internal validation than the external one, which suggests that our model can personalize the global model

well. These empirical results show the robustness and success of federated personalization of CD²-pFed on medical images, in addition to natural image classification.

Table 4. Results of real-world medical images, *i.e.*, stain variant histology slides, on each clients. Performance are evaluated by client side test accuracy. The best results are marked in **bold**.

Methods	client #1	client #2	client #3	External (\uparrow)
FedAvg [28]	65.23	65.31	65.45	60.03
Local Train	75.53	74.87	74.21	34.31
LG-Fed [22]	76.32	76.90	77.01	63.22
FedPer [2]	75.43	75.21	75.56	57.89
Ours	77.39	77.45	77.38	65.66

Discussion. Comprehensive experiments on four datasets, characterized with different non-IID settings, confirmed that our CD²-pFed is the only method to consistently achieve state-of-the-art results. Although, LG-Fed performs better to FedAvg in the existence of feature skew, and FedPer better with a label distribution skew, their performance sharply declines when the non-IID settings are interchanged. It is assumed that LG-Fed personalizes the bottom layer to better learn from highly heterogeneous images, while FedPer personalizes the top layer to distinguish unbalanced samples. Our CD²-pFed, containing both the low-level and high-level personalization can reduce the reliance on prior knowledge for personalization decisions.

4.5. Ablation Analysis

The proposed CD²-pFed is composed of three functional components to assist the channel decoupling, namely the progressive personalization ratio increment scheme (LI), temporal average moving for the personalized weights (TA), and cyclic distillation (CD). To test the effectiveness of each scheme, we perform ablation studies on CIFAR-10 with a $s = 2$ split. As shown in Table 5, we can observe that 1) with all components, CD²-pFed achieves the best performance, demonstrating the effectiveness of integrating three schemes *i.e.* LI + TA + CD, with a 1.51% local test accuracy improvement to the origin channel decoupling; 2) CD achieves the highest improvement, TA second, and LI the least; 3) LI and TA can stabilize the training, resulting to a smaller stand deviation. We also visualize the training performance of non-IID CIFAR-10 in Figure 6, where our CD²-pFed achieves a faster convergence compared with existing layer-wise personalization methods, which requires fewer communication rounds during the federation.

5. Limitations and Conclusions

In this paper, we propose CD²-pFed to vertically decouple channels in the global model for personalization. Our vertical decoupling method can personalize the local model,

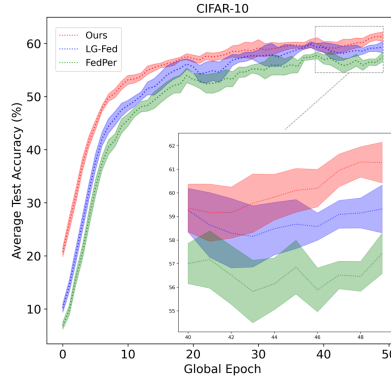


Figure 6. Effect on the performance of LeNet-5 on CIFAR-10, compared with different PFL frameworks [2, 22]. As illustrated, our CD²-pFed achieve higher test accuracy, together with a significant faster convergence speed.

Table 5. Ablation study on non-IID CIFAR-10 split, with $s = 2$, to evaluate the effectiveness of each component.

LI	TA	CD	Local (\uparrow)	New (\uparrow)
			90.31±0.67	59.12±0.32
✓			90.36±0.65	59.14±0.30
	✓		90.45±0.21	59.45±0.54
		✓	90.58±0.44	60.57±0.32
	✓	✓	91.67±0.54	61.20±2.03
✓		✓	91.00±1.03	59.84±1.53
✓	✓		90.81±0.38	59.34±0.56
✓	✓	✓	91.82±0.43	61.31±1.53

with guidance towards learning on high- and low-level feature representation. Subsequently, it can handle a variety of settings of data heterogeneity including the feature skew, label distribution skew, and concept shift. Empirically, compared with the previous layer-wise split which only learns one part of them, our framework shows a consistent success on four benchmark datasets. We also propose cyclic distillation to impose a consistency regularization and prevent the weights divergence in personalization. However, the cyclic distillation is currently designed by using soft predictions, which restricted itself to classification tasks. The extensions to segmentation and detection are left to future work. To stabilize the training process, we leverage a temporal average moving for personalized weights and a progressive increase scheme for the personalization ratio. Yet, we primarily assign a fixed personalization ratio for all layers, which yields an interesting future direction on searching a layer-specific optimal ratio.

Acknowledgements. The work described in this paper is supported by grants from HKU Startup Fund and HKU Seed Fund for Basic Research (Project No. 202009185079).

References

- [1] George J Annas. Hipaa regulations—a new era of medical-record privacy?, 2003. **1**
- [2] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019. **1, 2, 3, 4, 5, 6, 7, 8**
- [3] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kidon, Jakub Konečný, Stefano Mazzocchi, H Brendan McMahan, et al. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019. **1**
- [4] Xuan Gong, Abhishek Sharma, Srikrishna Karanam, Ziyang Wu, Terrence Chen, David Doermann, and Arun Innanje. Ensemble attention distillation for privacy-preserving federated learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15076–15086, October 2021. **1**
- [5] Pengfei Guo, Puyang Wang, Jinyuan Zhou, Shanshan Jiang, and Vishal M. Patel. Multi-institutional collaborations for improving deep learning-based magnetic resonance image reconstruction using federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2423–2432, June 2021. **1**
- [6] Filip Hanzely and Peter Richtárik. Federated learning of a mixture of global and local models. *arXiv preprint arXiv:2002.05516*, 2020. **1, 2**
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. **1**
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. **6**
- [9] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. **3**
- [10] Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019. **2, 3**
- [11] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019. **2, 3, 4**
- [12] Jing Ke, Yiqing Shen, and Yizhou Lu. Style normalization in histology with federated learning. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 953–956. IEEE, 2021. **6**
- [13] Jing Ke, Yiqing Shen, Jason D Wright, Naifeng Jing, Xiaoyao Liang, and Dinggang Shen. Identifying patch-level msi from histological images of colorectal cancer by a knowledge distillation model. In *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1043–1046. IEEE, 2020. **3**
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. **5**
- [15] Viraj Kulkarni, Milind Kulkarni, and Aniruddha Pant. Survey of personalization techniques for federated learning. In *2020 Fourth World Conference on Smart Trends in Systems, Security and Sustainability (WorldS4)*, pages 794–797. IEEE, 2020. **3**
- [16] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. **5**
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. **6**
- [18] Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581*, 2019. **4**
- [19] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10713–10722, June 2021. **1**
- [20] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018. **3**
- [21] Xiaoxiao Li, Meirui JIANG, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. In *International Conference on Learning Representations*, 2020. **3**
- [22] Paul Pu Liang, Terrance Liu, Liu Ziyin, Nicholas B Allen, Randy P Auerbach, David Brent, Ruslan Salakhutdinov, and Louis-Philippe Morency. Think locally, act globally: Federated learning with local and global representations. *arXiv preprint arXiv:2001.01523*, 2020. **1, 2, 3, 4, 5, 6, 7, 8**
- [23] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *arXiv preprint arXiv:2006.07242*, 2020. **4**
- [24] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. Feddg: Federated domain generalization on medical image segmentation via episodic learning in continuous frequency space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1013–1023, June 2021. **1**
- [25] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015. **1**
- [26] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. **4**
- [27] Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*, 2020. **3**
- [28] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. **1, 2, 4, 5, 6, 7, 8**

- [29] Joaquin Quiñero-Candela, Masashi Sugiyama, Neil D Lawrence, and Anton Schwaighofer. *Dataset shift in machine learning*. Mit Press, 2009. 2, 3
- [30] Jian Ren, Xiaohui Shen, Zhe Lin, Radomir Mech, and David J Foran. Personalized image aesthetics. In *Proceedings of the IEEE international conference on computer vision*, pages 638–647, 2017. 3, 6
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28:91–99, 2015. 1
- [32] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 3
- [33] Johannes Schneider and Michail Vlachos. Personalization of deep learning. *arXiv preprint arXiv:1909.02803*, 2019. 3
- [34] Leslie N Smith. Cyclical learning rates for training neural networks. In *2017 IEEE winter conference on applications of computer vision (WACV)*, pages 464–472. IEEE, 2017. 4
- [35] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. Federated multi-task learning. *arXiv preprint arXiv:1705.10467*, 2017. 2, 3
- [36] Jingwei Sun, Ang Li, Binghui Wang, Huanrui Yang, Hai Li, and Yiran Chen. Soteria: Provable defense against privacy leakage in federated learning from representation perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9311–9319, June 2021. 1
- [37] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *arXiv preprint arXiv:2103.00710*, 2021. 2
- [38] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *arXiv preprint arXiv:1703.01780*, 2017. 3
- [39] Paul Voigt and Axel Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, 10:3152676, 2017. 1
- [40] Qiong Wu, Kaiwen He, and Xu Chen. Personalized federated learning for intelligent iot applications: A cloud-edge based framework. *IEEE Open Journal of the Computer Society*, 1:35–44, 2020. 3
- [41] Guodong Xu, Ziwei Liu, Xiaoxiao Li, and Chen Change Loy. Knowledge distillation meets self-supervision. In *European Conference on Computer Vision*, pages 588–604. Springer, 2020. 3
- [42] Junho Yim, Donggyu Joo, Jihoon Bae, and Junmo Kim. A gift from knowledge distillation: Fast optimization, network minimization and transfer learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4133–4141, 2017. 3
- [43] Jiahui Yu and Thomas S Huang. Universally slimmable networks and improved training techniques. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1803–1811, 2019. 5
- [44] Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging federated learning by local adaptation. *arXiv preprint arXiv:2002.04758*, 2020. 3
- [45] Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging federated learning by local adaptation. *arXiv preprint arXiv:2002.04758*, 2020. 3
- [46] Sukmin Yun, Jongjin Park, Kimin Lee, and Jinwoo Shin. Regularizing class-wise predictions via self-knowledge distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13876–13885, 2020. 3
- [47] Lin Zhang, Yong Luo, Yan Bai, Bo Du, and Ling-Yu Duan. Federated learning for non-iid data via unified feature learning and optimization objective alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4420–4428, October 2021. 1
- [48] Zizhao Zhang, Han Zhang, Sercan O Arik, Honglak Lee, and Tomas Pfister. Distilling effective supervision from severe label noise. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9294–9303, 2020. 3
- [49] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018. 2, 5
- [50] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. Collaborative unsupervised visual representation learning from decentralized data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4912–4921, October 2021. 1