

PGFed: Personalize Each Client’s Global Objective for Federated Learning

Jun Luo

Intelligent Systems Program,
University of Pittsburgh, USA

jul117@pitt.edu

Matias Mendieta, Chen Chen

Center for Research in Computer Vision
University of Central Florida, USA

mendieta@knights.ucf.edu; chen.chen@crvc.ucf.edu

Shandong Wu

Department of Radiology, Intelligent Systems Program,
Department of Biomedical Informatics, Department of Bioengineering,
University of Pittsburgh, USA

wus3@upmc.edu

Abstract

The mediocre performance of conventional federated learning (FL) over heterogeneous data has been facilitating personalized FL solutions, where, unlike conventional FL which trains a single global consensus model, different models are allowed for different clients. However, in most existing personalized FL algorithms, the collaborative knowledge across the federation was only *implicitly* passed to the clients in ways such as model aggregation or regularization. We observed that this implicit knowledge transfer fails to maximize the potential value of each client’s empirical risk toward other clients. Based on our observation, in this work, we propose **Personalized Global Federated Learning (PGFed)**, a novel personalized FL framework that enables each client to **personalize** its own **global** objective by **explicitly** and adaptively aggregating the empirical risks of itself and other clients. To avoid massive ($O(N^2)$) communication overhead and potential privacy leakage, each client’s risk is estimated through a first-order approximation for other clients’ adaptive risk aggregation. On top of PGFed, we develop a momentum upgrade, dubbed **PGFedMo**, to more efficiently utilize clients’ empirical risks. Our extensive experiments under different federated settings with benchmark datasets show consistent improvements of PGFed over the compared state-of-the-art alternatives.

1. Introduction

Training deep learning models requires big data. Yet, datasets often reside on parties (devices or silos) with strict




Figure 1. In PGFed, the explicit collaborative knowledge transfer shows in the design of the local objective as a “personalized global” objective at client i , where non-local risks are involved.

privacy rules [37], making it infeasible to migrate and gather the data for centralized training. Over the past years, federated learning [14, 20, 26, 38] (FL) has shown great potential in training deep learning models over decentralized parties, avoiding the need to gather the private data.

Conventional FL aims to train a single global consensus model by orchestrating the participating clients with a central server. The most notable FL algorithm, FedAvg [26], proceeds the training by communicating and exchanging only the locally trained and server aggregated global model between clients and server, leaving the private datasets intact. As a privacy-preserving machine learning technique, FL tremendously boosts new collaborations of decentralized parties in a number of areas [14, 16, 23, 40].

Unfortunately, the utility of FedAvg over heterogeneous (non-IID) data was questioned in recent years [13, 21, 30, 31]. Specifically, FedAvg suffers from non-guaranteed convergence over heterogeneous data [30] and poor generalizability after fine-tuning [13]. This promotes numerous solutions that try to improve on top of FedAvg by further

restricting the learning of the global model [1, 15, 19, 21].

Meanwhile, another direction of efforts trying to mitigate the data heterogeneity challenge shifts gears towards personalized FL [18, 35], where different models are allowed for different clients. Some efforts in this direction focus on different personalized layers and optimization techniques [5, 22, 29]. Leveraging multi-task [44] or meta-learning [10] is also shown to be beneficial in personalized FL [3, 7, 33]. Other works include Clustered FL [9, 32], interpolation of personalized models [6, 24, 43], etc [25, 39, 42].

However, in most existing personalized FL algorithms [2, 4–7, 12, 22, 28], *the way in which the collaborative knowledge passes down from the server to the clients is implicit*. Here, we consider the collaborative knowledge as non-local information, such as the global objective of FedAvg, $F(\theta) = \sum_i p_i F_i(\theta)$, where θ is the global model and $F_i(\cdot)$ represents client i 's local objective whose weights are denoted as p_i 's. In addition, we define “*implicitness*” by defining its opposite side, “*explicitness*”, as a direct engagement with multiple clients’ (empirical) risks. For instance, updating the global model of FedAvg is explicit, where the direct engagement is achieved through communication. However, this can hardly be the case for updating clients’ personalized models, as it would take $O(N^2)$ communication overhead to transmit each client’s personalized model to every client, assuming FedAvg’s communication cost is $O(N)$ over N clients. Consequently, most personalized FL algorithms implicitly pass the collaborative knowledge from the server to the clients by embedding it into the aggregation of model weights or as different kinds of regularizers.

Why should we care about the “*explicitness*”, especially for updating the personalized models? Let us first assume all communication has zero cost, and, as an example, design a client’s explicit local objective as a “*personalized global objective*” in the same form as the global objective of FedAvg (weighted sum of all clients’ risks), i.e. $F_i(\theta_i) = f_i(\theta_i) + \mu/(N - 1) \sum_{j \neq i} f_j(\theta_j)$, where $F_i(\cdot)$ and $f_i(\cdot)$ are client i 's local objective and empirical risk, respectively, and μ is a hyperparameter.

Although in a simple form, such an exemplar design facilitates the generalizability of the personalized models directly by penalizing its performance over other clients’ risks, contributing to better local performance. We demonstrate this characteristic by an empirical study: we personalize the output global model of FedAvg through S steps of local gradient-based update, supposing that the $O(SN^2)$ communication cost is affordable for now. We compare the exemplar explicit design of each client’s local objective with a simple implicit way – each client’s objective is only its own empirical risk.

The result of this empirical study are shown in Fig. 2,

on CIFAR10 dataset with 100 heterogeneous clients, we observe that the explicit transfer of collaborative knowledge exhibits a stronger ability to adapt the model towards clients’ local data than the implicit counterpart, with a mean individual performance gain of 7.73% over local test data from the initial global model, 2.33% higher than that of an implicit personalization. *Intriguing as the results may seem, how can we transfer this idea back to the communication-expensive real-world FL settings where acquiring all $f_j(\theta_j) \forall i, j \in [N]$ will cost massive communication overhead?* Our solution is to estimate $f_j(\theta_j)$ by approximation as shown in Fig. 1.

Based on the above observation, in this work, we propose **Personalized Global Federated Learning (PGFed)**, a novel personalized FL framework that enables each client to **personalize** its own **global** objective by *explicitly and adaptively* aggregating the empirical risks of itself and other clients. To avoid massive communication overhead and potential privacy leakage, each client’s risk is estimated through a first-order approximation for other clients’ adaptive risk aggregation. Thereby, the clients are able to explicitly acquire the collaborative knowledge, and their personalized models can enjoy better generalizability. We summarize our contribution as follows:

- We uncover that the explicitness of a personalized FL algorithm empowers itself with stronger adaptation ability. Based on this observation, we propose PGFed, a novel explicit personalized FL algorithm that frames the local objective of each client as a personalized global objective.
- To the best of our knowledge, PGFed is the first work in the field that explicitly transfers the global collaborative knowledge from the server to the clients, and at the same time circumvents massive communication costs by introducing estimates of non-local empirical risks.
- On top of PGFed, we develop a momentum upgrade, dubbed PGFedMo, to more efficiently utilize clients’ empirical risks.
- We evaluate PGFed and PGFedMo on benchmark datasets under different FL settings. The results show that both algorithms outperform the compared state-of-the-art personalized FL methods.

2. Related Work

Federated Learning. Federated learning [14, 20, 38] focuses on training a global consensus model over a federation of clients with similar data. The most notable FL algorithm, FedAvg [26], proceeds the training in federated rounds iteratively, where copies of the global model are broadcast to selected clients. Locally updated versions of the copies are then aggregated as the new global model for the next round.

However, evidence began to accumulate in recent years that FedAvg is vulnerable towards data heterogeneity

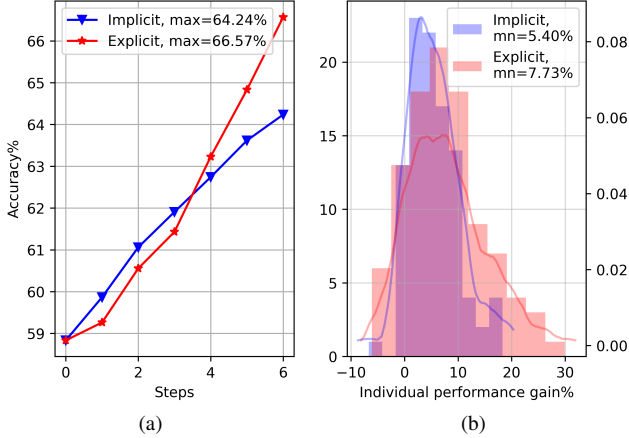


Figure 2. On the task of personalizing the output of FedAvg on CIFAR10 with 100 heterogeneous clients, the performance of an exemplar explicit objective ($F_i(\theta_i) = f_i(\theta_i) + \mu/(N-1) \sum_{j \neq i} f_j(\theta_i)$) and a simple implicit objective ($F_i(\theta_i) = f_i(\theta_i)$), assuming the communication cost is affordable. Fig. 2a shows the trend of the mean personalized test accuracy. Fig. 2b shows the histogram and density estimation of the final individual gain.

(non-IID data) [13, 21, 30, 31], which is the most common data distribution scenario in real-world FL. This vulnerability on non-IID data manifests itself as non-guaranteed convergence [30] and poor generalizability on test data after fine-tuning [13], which can cost the incentives of participating clients.

To improve the performance of the global model, FedProx [21] adds a proximal term to clients’ empirical risks to restrict the local update. FedDyn [1] introduces a dynamic regularizer to each client to align the global and local optima. SCAFFOLD [15] corrects the local update by a control variate. And FedAlign [27] focuses on local model generality rather than proximal restrictions. However, while effective, these global FL algorithms cannot systematically alleviate the data heterogeneity challenge.

Personalized Federated Learning. Personalized FL [18, 35] allows different models for different clients. This relaxation fundamentally mitigates the impact caused by the divergence of local and global optima. Methods in this category are the ones we mainly compare with in our experiments.

In this branch of work, [2, 5, 22] focus on aggregating different layers of the model. A recent work, FedBABU [28], aggregates the feature extractor for a global model, and fine-tunes it for personalized models. Moreover, some works leverage multi-task [44] or meta-learning [10] to learn the relationships between clients’ data distributions [3, 7, 33]. For instance, Per-FedAvg [7] takes advantage of Model-

Agnostic Meta-Learning (MAML) [8] and trains an easy-to-adapt initial shared model. Instead of training a different model for each client, clustered FL [9, 32] trains a distinct model for each cluster of clients, where clients in the same cluster share some quantifiable similarities. Interpolation of personalized models is also a popular direction [6, 24, 43], where works concentrate more on a personalized way to aggregate the models from different clients.

In addition, some recent works [24, 43] afford to pay the $O(N^2)$ communication cost to transmit the clients’ models to other clients, but their massive communication cost only benefits the model aggregation, instead of paying more attention on aggregating clients’ risks, which makes their algorithms fall into the implicit category. Another recent work [4] bridges the personalized FL and global FL by training a global model with balanced risk and personalized adaptive predictors with empirical risk. However, it is still implicit, as the clients do not engage with others’ risks. In our work, we focus on an explicit way of personalization by introducing estimates of non-local empirical risks to each client without massive communication costs.

3. Problem Formulation

In this section, we formalize the problem of conventional and personalized FL, and further elaborate on the explicitness of personalized FL.

3.1. Conventional and Personalized FL

Conventional federated learning aims to train a global consensus model for a federation of clients with similar data. Take the most notable FL algorithm, FedAvg [26], as an example. For a federation of N clients, the goal of FedAvg is to minimize the global objective defined as:

$$\min_{\theta} F(\theta) = \sum_{i=1}^N p_i F_i(\theta), \quad (1)$$

where θ denotes the global model, $F_i(\cdot)$ represents the local objective of client i , and the weight p_i is often set as $p_i = n_i/n$ with $n = \sum_k n_k$ where n_k denotes the number of data samples on client k . In FedAvg, the local objective $F_i(\cdot)$ often measures client i ’s empirical risk, i.e.

$$F_i(\theta) = \mathbb{E}_{\xi \sim \mathcal{D}_i} f_i(\theta|\xi) \approx \sum_{k=1}^{n_i} f_i(\theta|\xi_k), \quad (2)$$

where \mathcal{D}_i represents the data distribution, ξ_k is the k -th data sample on client i , and we use the empirical risk $f_i(\cdot)$ to approximate the true risk on client i . To simplify the notations, in the rest of the paper, we drop the summation denote $f_i(\cdot)$ directly as $\sum_{k=1}^{n_i} f_i(\cdot|\xi_k)$ unless further clarified.

Personalized FL relaxes the number of models. Compared to conventional FL, each client i is allowed to have

its own personalized model θ_i , i.e. the goal of personalized FL is defined as:

$$\min_{\Theta} F(\Theta) = \min_{\theta_1, \dots, \theta_N} \sum_{i=1}^N p_i F_i(\theta_i), \quad (3)$$

where Θ is a $d \times N$ matrix with d being the number of dimensions of the model.

In addition, the number of participating clients (N) in FL is often large [14], and not all clients are able to participate in each federated round. Therefore, a subset of clients \mathcal{S} with $|\mathcal{S}| = M$ is selected for each round of training.

3.2. Implicit vs. Explicit Local Objective

As discussed in Sec. 1, we defined two types of ways of passing the collaborative knowledge from the server to the clients for personalized FL algorithms. The explicit way updates the model with a direct engagement with multiple clients' empirical risks. For instance, updating the global model of FedAvg achieves explicitness through communication. And in the empirical study (Fig. 2) in Sec. 1, we provided an example of explicit local objective as:

$$F_i(\theta_i) = f_i(\theta_i) + \frac{\mu}{(N-1)} \sum_{j \neq i} f_j(\theta_j). \quad (4)$$

It is obvious that engaging with all $f_j(\theta_j) \forall i, j \in [N]$ to update the clients' personalized models would take $O(N^2)$ communication overhead. Therefore, most personalized FL algorithms choose to implicitly pass the collaborative knowledge by embedding it into the aggregation of model weights or as different regularizers.

One pitfall is that, to see whether a personalized FL algorithm manages to afford the $O(N^2)$ communication overhead is *NOT* the criterion of checking its explicitness. [24, 43] are two recent works that did afford the communication cost on small federations, but since their personalized model updates do not involve other clients' risks, both algorithms still fall into the implicit category.

Note that the explicitness of a personalized FL algorithm does not prevent it from simultaneously possessing characteristics from the implicit counterpart. For example, in a personalized FL algorithm, taking the globally aggregated model as a round-beginning initialization for the clients' personalized model is an implicit characteristic that can co-exist with using Eq. (4) as an explicit personalization.

4. Method

In this section, we introduce the proposed algorithm, PGFed, in details. We clarify our global and local objectives, explain how we circumvent the massive communication cost, and introduce a momentum version of the algorithm, dubbed PGFedMo. The full procedure of PGFed is provided in Algorithm 1.

4.1. Objectives of PGFed

We adopt Eq. (3) as the global objective of PGFed. As mentioned in Sec. 3.2, an algorithm can simultaneously leverage characteristics from both explicit and implicit collaborative knowledge transfer. In each training round, we first implicitly pass the collaborative knowledge to the clients as the globally aggregated model, θ_{glob} , which serves as the round-beginning initialization for the clients' personalized models.

To address the explicitness, we design the local objective as a *personalized global objective* in the same form as the global objective of FedAvg (weighted sum of every clients' risks). Different from the previous exemplar design shown in Eq. (4), in PGFed, the local objective is defined as:

$$F_i(\theta_i, \alpha_i) = f_i(\theta_i) + \mu \sum_{j \in [N]} \alpha_{ij} f_j(\theta_j), \quad (5)$$

where α_i is a N -dimensional vector of **learnable** scalars $\alpha_{ij} > 0$ which denotes the personalized weights of $f_j(\cdot)$ on client i . In this way, each client is able to personalize how much other clients' risks should contribute to their own objective. Since M clients are selected for every round, we initialize every as $\alpha_{ij} = 1/M \forall i, j \in [N]$. This new local objective slightly changes the global objective to:

$$\min_{\Theta, \mathbf{A}} F(\Theta, \mathbf{A}) = \min_{\theta_1, \dots, \theta_N, \alpha_1, \dots, \alpha_N} \sum_{i=1}^N p_i F_i(\theta_i, \alpha_i), \quad (6)$$

where \mathbf{A} is an $N \times N$ matrix with the i -th row being α_i .

Now, the question becomes: how can we overcome the massive communication cost to achieve explicit personalization? We answer this question in the next subsection.

4.2. Non-local Risk Estimation

In most real-world FL settings, $O(N^2)$ communication cost is often not affordable. In PGFed, to achieve the explicitness without having to pay the unaffordable cost, we estimate the non-local empirical risks by first-order approximations. Specifically, we define the non-local empirical risk on client i as $f_j(\theta_i) \forall j \neq i$, and estimate this term by using Taylor expansion at θ_j , i.e.

$$f_j(\theta_i) \approx f_j(\theta_j) + \nabla_{\theta_j} f_j(\theta_j)^T (\theta_i - \theta_j), \quad (7)$$

where the higher-order terms are ignored. By plugging the approximation into Eq. (5), we have:

$$F_i(\theta_i, \alpha_i) \approx f_i(\theta_i) + \mathcal{R}_{aug}^{[N]}(\theta_i, \alpha_i), \quad (8)$$

$$\mathcal{R}_{aug}^{[N]}(\theta_i, \alpha_i) = \mu \sum_{j \in [N]} \alpha_{ij} (f_j(\theta_j) + \nabla_{\theta_j} f_j(\theta_j)^T (\theta_i - \theta_j)). \quad (9)$$

where we define $\mathcal{R}_{aug}^{[N]}(\cdot)$ as an *auxiliary risk* over all client $j \in [N]$. It is the existence of the auxiliary risk $\mathcal{R}_{aug}^{[N]}(\cdot)$ that makes the proposed algorithm an explicit personalization.

4.3. Gradient-based Update

In PGFed, the personalized models are updated through gradient-based optimizers such as stochastic gradient descent (SGD). In this section, we derive the gradient of the local objective of client i with respect to the personalized model θ_i and the personalized weights α_i for the non-local risks. Note that the gradient of the local objective is just the gradient of the local empirical risk plus the gradient of the auxiliary risk. Therefore, for the gradient w.r.t. θ_i , we have:

$$\begin{aligned} \nabla_{\theta_i} F_i(\theta_i, \alpha_i) &= \nabla_{\theta_i} f_i(\theta_i) + \nabla_{\theta_i} \mathcal{R}_{aug}^{[N]}(\theta_i, \alpha_i) \\ &= \nabla_{\theta_i} f_i(\theta_i) + \underbrace{\mu \sum_{j \in [N]} \alpha_{ij} \nabla_{\theta_j} f_j(\theta_j)}_{\tilde{g}_{[N]}}. \end{aligned} \quad (10)$$

Since the $\tilde{g}_{[N]}$, which we name the auxiliary gradient, is not related to θ_i , we can have it computed by the server and transmit it to client i , by letting client i upload α_i and client j upload $\nabla_{\theta_j} f_j(\theta_j)$. Similar for the weight vector α_i , we have:

$$\begin{aligned} \nabla_{\alpha_{ij}} F_i(\theta_i, \alpha_i) &= \mu (f_j(\theta_j) + \nabla_{\theta_j} f_j(\theta_j)^T (\theta_i - \theta_j)) \\ &= \underbrace{\mu (f_j(\theta_j) - \nabla_{\theta_j} f_j(\theta_j)^T \theta_j)}_{g_{\alpha}^{(1)}} + \underbrace{\mu \nabla_{\theta_j} f_j(\theta_j)^T \theta_i}_{g_{\alpha}^{(2)}}. \end{aligned} \quad (11)$$

Note that this gradient can be split into two components: the first term, $g_{\alpha}^{(1)}$, is a scalar purely associated with client j , which can be uploaded from client j with little cost; the second term is an interaction between the gradient of client j and the personalized model of client i . To accurately acquire $g_{\alpha}^{(2)}$ term itself would, again, cost $O(N^2)$ communication overhead. Therefore, we approximate this term by an average, $\bar{g}_{[N]}$, computed by the server, i.e.

$$g_{\alpha}^{(2)} \approx \bar{g}_{[N]}^T \theta_i = \frac{\mu}{N} \left(\sum_{j \in [N]} \nabla_{\theta_j} f_j(\theta_j) \right)^T \theta_i. \quad (12)$$

Although transmitting $\bar{g}_{[N]}$ can be costly, and it is possible for the server to transmit the scalar value of $g_{\alpha}^{(2)} \approx \bar{g}_{[N]}^T \theta_i$, we argue that it is not ideal to directly send this scalar to client i , since the calculations of $g_{\alpha}^{(2)}$ and $\nabla_{\alpha_{ij}} F_i(\theta_i, \alpha_i)$ are under the process of updating θ_i . Since calculating $g_{\alpha}^{(2)}$ involves θ_i , it would be more reasonable to treat $\bar{g}_{[N]}^T \theta_i$ as a variable and have it computed by client i locally, rather than treating it as a constant computed by the server.

In the descriptions of the proposed algorithm above, we did not elaborate on how the client selection procedure would affect the auxiliary risks of each client. As mentioned in Sec. 4.1, in each round t , a subset of clients, \mathcal{S}_t

Algorithm 1 PGFed and PGFedMo

Input: N clients, learning rates η_1, η_2 , number of rounds T , coefficient μ , momentum β for PGFedMo

Output: Personalized models $\theta_1^T, \dots, \theta_N^T$.

ServerExecute:

```

1: Initialize  $\alpha_{ij} \leftarrow 1/M \forall i, j \in [N]$ , global model  $\theta_{glob}^0$ 
2:  $\mathbf{A}[i] \leftarrow \alpha_i \forall i \in [N]$ 
3: for  $t \leftarrow 1, 2, \dots, T$  do
4:   Select a subset of  $M$  clients,  $\mathcal{S}_t$ 
5:    $g_t^{(1)} \leftarrow \{\}; \nabla_t \leftarrow \{\}$  // build for next round
6:   for  $i \in \mathcal{S}_t$  in parallel do
7:     if  $t=1$  then
8:        $\theta_i^t, g_{\alpha}^{(1)}, \nabla f(\theta_i^t), \alpha_i \leftarrow \text{ClientUpdate}(\theta_{glob}^{t-1}, t)$ 
9:     else
10:       $\tilde{g}_{\mathcal{S}_t} \leftarrow \mu \sum_{j \in \mathcal{S}_t} \alpha_{ij} \nabla_{t-1}[j]$ 
11:       $\bar{g}_{\mathcal{S}_t} \leftarrow \frac{\mu}{M} \left( \sum_{j \in \mathcal{S}_t} \nabla_{t-1}[j] \right)$ 
12:       $\theta_i^t, g_{\alpha}^{(1)}, \nabla f(\theta_i^t), \alpha_i \leftarrow \text{ClientUpdate}(\theta_{glob}^{t-1}, t,$ 
13:         $\tilde{g}_{\mathcal{S}_t}, \bar{g}_{\mathcal{S}_t}, g_{t-1}^{(1)})$ 
14:      end if
15:      // the next line records the values for next round
16:       $\mathbf{A}[i] \leftarrow \alpha_i; g_t^{(1)}[i] \leftarrow g_{\alpha}^{(1)}; \nabla_t[i] \leftarrow \nabla f(\theta_i^t)$ 
17:       $\theta_{glob}^t \leftarrow \sum_{i \in \mathcal{S}_t} p_i \theta_i^t$ 
18:    end for
19:    for  $i \in ([N] - \mathcal{S}_t)$  in parallel do
20:       $\theta_i^t \leftarrow \theta_i^{t-1}; \tilde{g}_i^t \leftarrow \tilde{g}_i^{t-1}$ 
21:    end for
22:  end for
23:  return  $\theta_1^T, \dots, \theta_N^T$ 

```

ClientUpdate($\theta_{global}^{t-1}, t, \tilde{g}, \bar{g}, g_{t-1}^{(1)})$:

```

1: if  $t=1$  then
2:    $\theta_i^t \leftarrow \text{ClientUpdate}(\theta_{global}^{t-1}, \eta_1)$  as in FedAvg
3: else
4:    $\theta_i^t \leftarrow \theta_{global}^{t-1}$ 
5:    $\tilde{g}_i^t \leftarrow \tilde{g}$  // without momentum
6:    $\tilde{g}_i^t \leftarrow (1 - \beta)\tilde{g} + \beta\tilde{g}_i^{t-1}$  // with momentum
7:   for Batch of data  $\mathcal{B} \in \mathcal{D}_i$  do
8:      $\theta_i^t \leftarrow \theta_i^t - \eta_1(\nabla f(\theta_i^t, \mathcal{B}) + \tilde{g}_i^t)$ 
9:      $g^{(2)} = \bar{g}^T \theta_i$ 
10:     $\forall j \in g_{t-1}^{(1)} : \alpha_{ij} \leftarrow \alpha_{ij} - \eta_2(g_{t-1}^{(1)}[j] + g^{(2)})$ 
11:   end for
12: end if
13:  $g_{\alpha}^{(1)} \leftarrow \mu (f(\theta_i^t) - \nabla f(\theta_i^t)^T \theta_i^t)$  // for next round
14: return  $\theta_i^t, g_{\alpha}^{(1)}, \nabla f(\theta_i^t), \alpha_i$ 

```

with $|\mathcal{S}_t| = M$, are selected for the training. We, therefore, slightly modify the auxiliary risk from $\mathcal{R}_{aug}^{[N]}$ to $\mathcal{R}_{aug}^{\mathcal{S}_t}$. This change will subsequently affect the scope of two terms, namely the auxiliary gradient $\tilde{g}_{[N]}$ w.r.t. θ_i (changed to

	CIFAR10			CIFAR100		
	25 clients	50 clients	100 clients	25 clients	50 clients	100 clients
Local	72.40 ± 0.45	70.28 ± 0.38	67.39 ± 0.20	32.74 ± 0.08	26.05 ± 0.34	23.06 ± 0.47
FedAvg [26]	65.07 ± 0.25	64.41 ± 0.66	63.19 ± 0.46	28.48 ± 0.59	26.06 ± 0.65	25.58 ± 0.80
FedDyn [1]	67.31 ± 0.36	65.02 ± 0.91	62.49 ± 0.06	34.17 ± 0.43	27.06 ± 0.18	23.88 ± 0.36
pFedMe [34]	70.60 ± 0.23	68.92 ± 0.35	66.40 ± 0.04	27.97 ± 0.24	23.82 ± 0.06	22.35 ± 0.03
FedFomo [43]	72.33 ± 0.03	72.17 ± 0.48	70.86 ± 0.27	32.15 ± 0.61	25.90 ± 1.17	24.48 ± 0.44
APFL [6]	77.03 ± 0.26	77.36 ± 0.18	76.29 ± 0.13	39.16 ± 0.93	35.15 ± 0.65	33.86 ± 0.60
FedRep [5]	76.85 ± 0.44	76.03 ± 0.17	72.30 ± 0.52	33.43 ± 0.80	26.86 ± 0.39	22.76 ± 0.45
LG-FedAvg [22]	72.83 ± 0.28	70.44 ± 0.31	67.55 ± 0.09	33.65 ± 0.19	27.13 ± 0.37	24.82 ± 0.28
FedPer [2]	77.84 ± 0.18	77.76 ± 0.22	75.01 ± 0.20	35.22 ± 0.67	28.63 ± 0.70	25.56 ± 0.26
Per-FedAvg [7]	75.80 ± 0.56	76.53 ± 0.31	77.13 ± 0.32	34.45 ± 0.49	33.57 ± 0.80	36.40 ± 0.52
FedRoD [4]	79.73 ± 0.68	79.61 ± 0.22	77.76 ± 0.32	39.55 ± 0.58	33.87 ± 2.42	31.49 ± 0.19
FedBABU [28]	78.92 ± 0.36	79.35 ± 0.84	76.34 ± 0.22	32.71 ± 0.23	29.66 ± 0.64	27.72 ± 0.11
PGFed (ours)	81.02 ± 0.41	81.42 ± 0.31	78.56 ± 0.35	43.12 ± 0.03	38.45 ± 0.44	35.71 ± 0.54
PGFedMo (ours)	81.20 ± 0.08	81.48 ± 0.32	78.74 ± 0.22	43.44 ± 0.14	38.50 ± 0.45	35.76 ± 0.65

Table 1. Mean top-1 personalized accuracy of the proposed algorithms and different baselines. For each FL setting and each method, we report the mean and standard deviation over three different seeds. The best performance is presented in **bold** font.

$\tilde{\mathbf{g}}_{S_t}$), and $\bar{\mathbf{g}}_{[N]}$, a portion of the gradient w.r.t. α_{ij} (changed to $\bar{\mathbf{g}}_{S_t}$). We formally define these two new terms below:

$$\tilde{\mathbf{g}}_{S_t} = \mu \sum_{j \in S_t} \alpha_{ij} \nabla_{\theta_j} f_j(\theta_j) \quad (13)$$

$$\bar{\mathbf{g}}_{S_t} = \frac{\mu}{M} \left(\sum_{j \in S_t} \nabla_{\theta_j} f_j(\theta_j) \right) \quad (14)$$

4.4. PGFed with Momentum (PGFedMo)

When considering the client selection for each round, the auxiliary risk changed from $\mathcal{R}_{aug}^{[N]}$ to $\mathcal{R}_{aug}^{S_t}$. This reduces the number of clients’ risks involved in the calculation of auxiliary risk from N to M , which could be huge. To compensate for this loss, each client maintains $\tilde{\mathbf{g}}_{S_t}^i$ locally, and every time when the client is selected, it updates its local auxiliary gradient $\tilde{\mathbf{g}}_{S_t}^i$ in a momentum manner, i.e.

$$\tilde{\mathbf{g}}_{S_t}^i = (1 - \beta) \tilde{\mathbf{g}}_{S_t}^i(\text{downloaded}) + \beta \tilde{\mathbf{g}}_{S_t}^i(\text{previous}) \quad (15)$$

In this way, the clients are able to carry the auxiliary gradient, without having to discard the collaborative knowledge from the $N - M$ clients each round. We summarize the proposed algorithm in Algorithm 1.

5. Experiments

5.1. Experimental Setup

Compared methods. We evaluate the proposed algorithm against a number of state-of-the-art FL algorithms, including two global FL algorithms: the leading FedAvg

and FedDyn that uses a dynamically regularized local training. Compared personalized FL algorithms include pFedME which conducts personalization with Moreau Envelopes; FedFomo and APFL which interpolates personalized models by interpolation; FedRep, LG-FedAvg, and FedPer which personalize different layers of the model; Per-FedAvg which leverages meta-learning to learn an initial shared model, and train one SGD step over test data; FedRoD which trains a global model with balanced risk and personalized adaptive predictors with empirical risk; and FedBABU which focuses on global representation learning and personalize the model by fine-tuning for one epoch.

Datasets. We conduct experiments on different benchmark datasets. We report the results on CIFAR10 and CIFAR100 [17], and leave the experiments on other datasets in the supplementary materials. Following [11], we use Dirichlet distribution with $\alpha = 0.3$ to partition the dataset into heterogeneous settings with 25, 50, and 100 clients. For each setting, each client’s local training and test datasets are under the same distribution. We report the mean top-1 personalized accuracy by a simple average of the local test accuracies over all clients.

Implementation details. We adopt the same convolutional neural network (CNN)’s architecture as in FedAvg for all the compared methods. There are 2 convolutional layers with 32, and 64 5×5 kernels and 2 fully connected layers with 512 hidden units in this architecture. For the methods that involve representation learning, we split the model into a feature extractor with the first three layers, and

a classifier with the last layer. We use a stochastic gradient descent (SGD) optimizer for every method with a fixed momentum of 0.9. Each method is trained on the dataset for 300 federated rounds. For each federated setting, the client sample rate is set to 25%, and the local training epochs is set to 5. The learning rate is tuned from $\{0.1, 0.01, 0.001, 0.0001\}$ for every method. For PGFed, the coefficient for the auxiliary risk, μ , is tuned from $\{0.1, 0.05, 0.01, 0.005, 0.001\}$. For PGFedMo, the momentum, β , is tuned from $\{0.2, 0.5, 0.8\}$. Hyperparameters values for the compared baselines are included in the supplementary materials.

5.2. Main Results and Analysis

We report the mean and standard deviation of three different seeds for each setting. As shown in Tab. 1, our method achieves the highest mean accuracies on all but one setting on CIFAR100. With large data heterogeneity using Dir(0.3) as in our experiments, local training can often achieve better performance than participating in global FL. As the number of clients increases, the clients benefit more from FL, since local training is more likely to overfit on the small local datasets. In personalized FL algorithms, PGFed outperforms the methods that personalize different components of the model such as FedRep, LG-FedAvg, and FedPer, and even with the fine-tuning in FedBABU, as well as the interpolation methods such as APFL and FedFomo. With the global model and the personalized adapter, FedRod benefits the local performance by the most within the baselines, yet was still exceeded by PGFed by up to 4.73% (on CIFAR100 with 50 clients). Per-FedAvg on CIFAR100 with 100 clients has the highest performance among the compared methods. Besides Per-FedAvg’s easy-to-adapt nature of its meta-learned shared model, the one step of training on the test set largely benefits the personalized model when the number of clients is large. For PGFed, adding momentum can further boost its performance since more clients’ risk is included in the clients’ auxiliary risks with the momentum update.

	25 clients		50 clients		100 clients	
	#round	speedup	#round	speedup	#round	speedup
APFL	31	1.0×	28	1.3×	24	2.2×
FedPer	8	3.9×	6	5.8×	8	6.5×
Per-FedAvg	30	1.0×	33	1.1×	52	1.0×
FedRoD	26	1.2×	35	1.0×	10	5.2×
PGFed	9	3.4×	14	2.5×	15	3.5×
PGFedMo	9	3.4×	14	2.5×	15	3.5×

Table 2. The number of rounds to achieve 70% mean top-1 personalized accuracy on CIFAR10. The speedup is computed based on the slowest approach listed (i.e. “1.0×”).

In Tab. 2, we show the convergence speed of PGFed and PGFedMo by reporting the number of rounds at which the algorithm achieves 70% mean top-1 personalized accuracy on CIFAR10. For each setting, we set the algorithm that takes the most round to reach 70% as “1.0×”, and find that the proposed PGFed and PGFedMo both consistently present a decent amount of speedup (3.1× on average) from the compared personalized FL algorithms. Among the six compared personalized FL algorithms, FedPer has the fastest convergence speed in every setting. The reason for the fast convergence could be the smaller size of the globally aggregated model component (only the feature extractor layers of a CNN model are aggregated in FedPer). Although PGFed and PGFedMo are not the fastest, the personalized models by the proposed algorithms converge with better generalizability than the faster FedPer (see Tab. 1).

	25 clients	50 clients	100 clients
FedAvg	-8.99±10.36	-8.90±15.48	-5.02±14.30
APFL	2.79 ± 8.07	5.73 ± 8.43	8.37 ± 6.91
FedPer	5.31 ± 2.56	8.31 ± 6.00	8.63 ± 5.26
Per-FedAvg	0.62 ± 8.72	4.83 ± 7.96	9.66 ± 8.73
FedRoD	7.80 ± 3.68	8.84 ± 6.29	10.68±6.14
PGFed	8.49 ± 4.67	10.78±5.88	11.15±5.06
PGFedMo	8.61 ± 3.59	10.90±6.11	11.16±5.44

Table 3. Mean and standard deviation of the individual performance gain over Local training in terms of accuracy% on local test set on CIFAR10.

	25 clients	50 clients	100 clients
FedAvg	-3.29 ± 4.22	0.02 ± 4.63	1.77 ± 6.38
APFL	6.48 ± 2.93	8.70 ± 3.37	9.31 ± 4.55
FedPer	3.43 ± 1.80	2.16 ± 2.45	2.31 ± 3.54
Per-FedAvg	1.05 ± 3.79	7.27 ± 4.64	12.27±5.70
FedRoD	7.32 ± 2.68	6.59 ± 3.17	7.47 ± 3.69
PGFed	9.34 ± 1.71	9.01 ± 2.97	12.05±3.93
PGFedMo	9.40 ± 1.87	8.99 ± 2.76	12.07±3.97

Table 4. Mean and standard deviation of the individual performance gain over Local training in terms of accuracy% on local test set on CIFAR100.

Besides the overall performance and the convergence speed, we take a micro-perspective to examine the individual performance gain over the local training strategy. Specifically, we concentrate on the statistics of the individual performance gain across the federation of clients, which can indicate the fairness performance of the algorithms. For instance, a high mean individual gain with a high standard deviation might indicate that the large per-

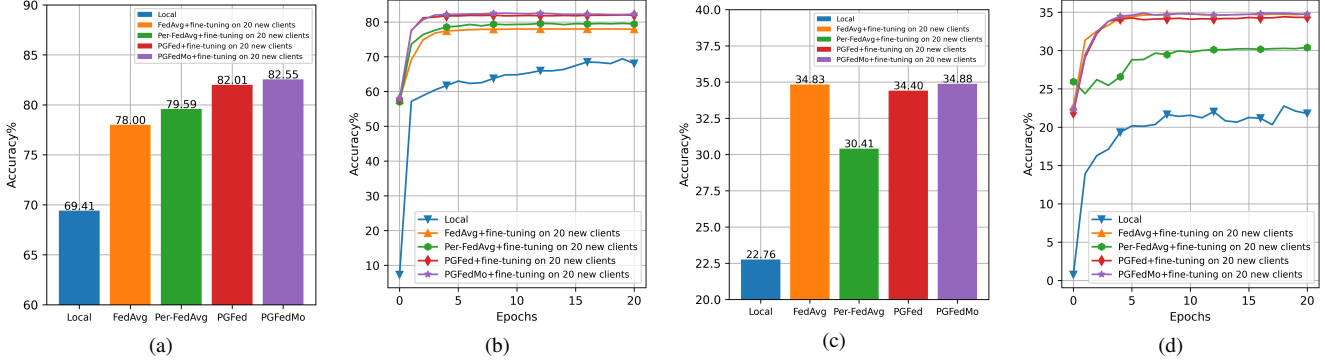


Figure 3. A comparison w.r.t. generalizability of different approaches on new clients on CIFAR10 and CIFAR100. We fine-tune the global models of different FL approaches and compare them with `Local` training. For FL approaches, we first train the global models on 80 clients for 150 rounds before fine-tuning them on 20 new clients for 20 epochs. Fig. 3a and Fig. 3c show the final mean personalized accuracies over the 20 new clients on CIFAR10 and CIFAR100, respectively. Fig. 3b and Fig. 3d show the learning curve along the fine-tuning.

formance gain on some clients come with a sacrifice over the performance gain on other clients. In Tab. 3 and 4, we show that while enjoying the highest mean individual gain in general, the standard deviations of both `PGFed` and `PGFedMo` are reasonably small as well, suggesting a fair boost of performance over all the clients. On CIFAR100 with 100 clients, `Per-FedAvg` has 0.2% higher mean individual gain than `PGFedMo`, but the standard deviation is 1.73% higher, which is not ideal. We attribute the high mean and small standard deviation merit of the proposed algorithms to the explicitness of our design of local objectives. Although clients’ local objectives are personalized, the explicitness enables the clients’ local objectives to be of the same form. And each client’s local empirical risk takes roughly $1/(1 + \mu)$ of its own objective, which is the same for every client, hence the similar individual performance gain and fair federated personalization.

5.3. Generalizability to New Clients

In real-world FL settings, it is possible that some clients did not participate in the FL training, but wish to have a model that could quickly adapt to their local data. It is especially desirable for a personalized FL algorithm to possess such an ability. By design, `PGFed` is able to generate a global model as a side product (line 16 of Algorithm 1). Therefore, we simulate such a setting and study the generalizability of our global model, and compare it against several baseline (global and personalized) algorithms that also produce the global model. We simulate this setting by conducting FL over 80 clients randomly selected from 100 clients on CIFAR10 and CIFAR100 datasets, and fine-tuning the trained global model on the rest 20 new clients. Besides `Local` training, we compare `PGFed` and `PGFedMo` with `FedAvg` and `Per-FedAvg`, whose goal is to train an easy-to-adapt model. The results are shown in Fig. 3.

Since the local fine-tuning on the 20 new clients is a standard vanilla training with SGD for all five compared methods, a high mean personalized accuracy directly indicates stronger overall generalizability of the global model (for FL algorithms), especially when the starting mean accuracies (before fine-tuning) are roughly the same. From Fig. 3, we can see that the `Local` trained models do not generalize well on local test data, since the size of the local dataset is likely to be small with under a 100-client setting. For the FL algorithms, `PGFedMo` achieves the best mean personalized accuracy, which shows the strong generalizability of the global model of the proposed algorithm thanks to the explicit design. While not the main focus, the strong generalizability of `PGFed`’s and `PGFedMo`’s global models shown on new clients also indicates strong adaptiveness in the original personalized FL task.

6. Conclusion and Discussion

In this work, we uncover that the explicitness of a personalized FL algorithm empowers itself with stronger adaptation ability. Based on our observations, we propose, `PGFed`, and its momentum upgrade, `PGFedMo`. Both algorithms explicitly transfer the global collaborative knowledge from the server to the clients, and at the same time circumvents massive communication costs by introducing estimates of non-local empirical risks. Our extensive experiments demonstrate the improvements of the proposed algorithms over state-of-the-art personalized alternatives on benchmark datasets under different heterogeneous FL settings.

We expect the proposed framework to be extended in different directions. First, we empirically show that explicit auxiliary risk aggregation can boost the personalized models’ generalizability, but whether it can reach a better

convergence rate in theory is still a question. In addition, since the proposed framework is agnostic, it can be easily combined with existing implicit FL algorithms such as the ones focusing on model weights aggregation, or different local regularizers. Another direction can be extended from the limitation of our work: although the proposed work manages to avoid the $O(N^2)$ communication overhead and achieves asymptotically the same communication overhead as FedAvg ($O(N)$), since each client is required to download 3 and upload 2 models/gradients per round (see Algorithm 1), on average the communication cost is still high (roughly 2.5 times as much as that of FedAvg). A more communication-efficient way to improve the local model’s generalizability is also worth investigations. We leave these directions for future work.

Acknowledgements

This work was supported in part by a National Institutes of Health (NIH) / National Cancer Institute (NCI) grant (1R01CA218405), a National Science Foundation (NSF) grant (CICI: SIVD: 2115082), the grant 1R01EB032896 as part of the NSF/NIH Smart Health and Biomedical Research in the Era of Artificial Intelligence and Advanced Data Science Program, a Pitt Momentum Funds scaling award (Pittsburgh Center for AI Innovation in Medical Imaging), and an Amazon AWS Machine Learning Research Award. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by NSF grant number ACI-1548562. Specifically, it used the Bridges-2 system, which is supported by NSF award number ACI-1928147, at the Pittsburgh Supercomputing Center. This work was also partially supported by the NSF/Intel Partnership on MLWiNS under Grant No. 2003198.

References

- [1] Durmus Alp Emre Acar, Yue Zhao, Ramon Matas, Matthew Mattina, Paul Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. In *International Conference on Learning Representations*, 2020. 2, 3, 6
- [2] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019. 2, 3, 6
- [3] Fei Chen, Mi Luo, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. Federated meta-learning with fast convergence and efficient communication. *arXiv preprint arXiv:1802.07876*, 2018. 2, 3
- [4] Hong-You Chen and Wei-Lun Chao. On bridging generic and personalized federated learning for image classification. In *International Conference on Learning Representations*, 2022. 2, 3, 6
- [5] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International Conference on Machine Learning*, pages 2089–2099. PMLR, 2021. 2, 3, 6
- [6] Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*, 2020. 2, 3, 6
- [7] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems*, 33:3557–3568, 2020. 2, 3, 6
- [8] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, pages 1126–1135. PMLR, 2017. 3
- [9] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020. 2, 3
- [10] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5149–5169, 2021. 2, 3
- [11] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019. 6
- [12] Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized cross-silo federated learning on non-iid data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 7865–7873, 2021. 2
- [13] Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*, 2019. 1, 3
- [14] 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. 1, 2, 4
- [15] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pages 5132–5143. PMLR, 2020. 2, 3
- [16] Latif U Khan, Walid Saad, Zhu Han, Ekram Hossain, and Choong Seon Hong. Federated learning for internet of things: Recent advances, taxonomy, and open challenges. *IEEE Communications Surveys & Tutorials*, 2021. 1
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [18] 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. 2, 3

- [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*, pages 10713–10722, 2021. 2
- [20] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020. 1, 2
- [21] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2:429–450, 2020. 1, 2, 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. 2, 3, 6
- [23] Guodong Long, Yue Tan, Jing Jiang, and Chengqi Zhang. Federated learning for open banking. In *Federated learning*, pages 240–254. Springer, 2020. 1
- [24] Jun Luo and Shandong Wu. Adapt to adaptation: Learning personalization for cross-silo federated learning. *arXiv preprint arXiv:2110.08394*, 2021. 2, 3, 4
- [25] 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. 2
- [26] 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, 3, 6
- [27] Matias Mendieta, Taojiannan Yang, Pu Wang, Minwoo Lee, Zhengming Ding, and Chen Chen. Local learning matters: Rethinking data heterogeneity in federated learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8397–8406, 2022. 3
- [28] Jaehoon Oh, Sangmook Kim, and Se-Young Yun. Fedbabu: Towards enhanced representation for federated image classification. *arXiv preprint arXiv:2106.06042*, 2021. 2, 3, 6
- [29] Krishna Pillutla, Kshitiz Malik, Abdel-Rahman Mohamed, Mike Rabbat, Maziar Sanjabi, and Lin Xiao. Federated learning with partial model personalization. In *International Conference on Machine Learning*, pages 17716–17758. PMLR, 2022. 2
- [30] Anit Kumar Sahu, Tian Li, Maziar Sanjabi, Manzil Zaheer, Ameet Talwalkar, and Virginia Smith. On the convergence of federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 3:3, 2018. 1, 3
- [31] Felix Sattler, Simon Wiedemann, Klaus-Robert Müller, and Wojciech Samek. Robust and communication-efficient federated learning from non-iid data. *IEEE transactions on neural networks and learning systems*, 31(9):3400–3413, 2019. 1, 3
- [32] Geet Shingi. A federated learning based approach for loan defaults prediction. In *2020 International Conference on Data Mining Workshops (ICDMW)*, pages 362–368. IEEE, 2020. 2, 3
- [33] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet Talwalkar. Federated multi-task learning. *arXiv preprint arXiv:1705.10467*, 2017. 2, 3
- [34] Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020. 6
- [35] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. 2, 3
- [36] Hemant Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5018–5027, 2017. 11, 13
- [37] 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):10–5555, 2017. 1
- [38] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, et al. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021. 1, 2
- [39] Kangkang Wang, Rajiv Mathews, Chloé Kiddon, Hubert Eichner, Françoise Beaufays, and Daniel Ramage. Federated evaluation of on-device personalization. *arXiv preprint arXiv:1910.10252*, 2019. 2
- [40] Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1):1–19, 2021. 1
- [41] Jiancheng Yang, Rui Shi, and Bingbing Ni. Medmnist classification decathlon: A lightweight automl benchmark for medical image analysis. In *IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 191–195, 2021. 11
- [42] Tao Yu, Eugene Bagdasaryan, and Vitaly Shmatikov. Salvaging federated learning by local adaptation. *arXiv preprint arXiv:2002.04758*, 2020. 2
- [43] Michael Zhang, Karan Sapra, Sanja Fidler, Serena Yeung, and Jose M Alvarez. Personalized federated learning with first order model optimization. *arXiv preprint arXiv:2012.08565*, 2020. 2, 3, 4, 6
- [44] Yu Zhang and Qiang Yang. A survey on multi-task learning. *arXiv preprint arXiv:1707.08114*, 2017. 2, 3

Appendix

A. Overview

We organize the supplementary material as follows:

- In Appendix B, we report additional results and analyses of PGFed, and conduct experiments on two other datasets, OrganAMNIST [41] & Office-home [36], with different FL settings.
- In Appendix C, we compare the local computational speed of the proposed algorithm, PGFed, with the baselines that achieved top performance in the experiments on CIFAR10 and CIFAR100.
- In Appendix D, we further propose PGFed-CE, a variation on top of PGFed to reduce both the communication and the computation cost simultaneously.
- In Appendix E, we report details in hyperparameters regarding our experiments.

B. Additional Experiments and Analyses

B.1. Convergence Behavior

We empirically study the convergence behavior of PGFed and the baselines that achieved high performance on CIFAR10 and CIFAR100. For each method, we plot its mean personalized test accuracy on CIFAR10 for the first 150 rounds of training under 25-, 50- and 100-client settings, as shown in Fig. 4.

From the results, we can see that, while achieving the highest accuracy performance, PGFed is also able to consistently converge faster than several of the baselines that reach high accuracies. Fast as it is under these settings, one limitation of PGFed is that the update of α_{ij} in PGFed only happens when client i and client j are selected in two consecutive rounds (client j is selected exactly one round before client i), which happens by chance. Therefore, this randomness might slightly limit the overall convergence behavior of PGFed, but it is the existence of α_{ij} 's that enables the personalization of how much each client values other clients' empirical risks, hence the higher accuracy.

B.2. Visualization of A

We study the matrix, A , of α_{ij} 's in PGFed (recall that α_{ij} is the coefficient of client j 's empirical risk on client i). Specifically, we visualize the **change** of all $\alpha_{ij} \forall i, j \in [N]$ on CIFAR10 with 25 clients. As a refresher, α_i , the i -th row of A , represents the vector of weights indicating how much client i values other clients' risks. The j -th column of A , A_j , represents the vector of weights indicating how much client j is valuable towards other clients' risks.

In theory (without the estimation, see Eq. (5) in the main paper), the gradient of α_{ij} equals $\mu f_j(\theta_i)$, which is non-negative. This suggests that α_{ij} will always converge to 0

given enough time, i.e. $\forall i, j \in [N], \lim_{t \rightarrow \infty} \alpha_{ij} = 0$. However, since $f_j(\theta_i)$ is estimated through Taylor expansion, and higher order terms are omitted, in practice of PGFed, α_{ij} has a chance to increase (pink and red blocks in Fig. 5a), or at least decreases more slowly (light blue in Fig. 5a), as long as client i 's local objective decreases. Therefore, whether α_{ij} changes in the positive or negative directions is a result of optimization, which depends on whether $f_j(\cdot)$ can be helpful to reduce the local objective of client i .

We quantify this "help" that client j can offer by $\overline{\Delta A_j}$: the average change of the j -th column of A (i.e. $\alpha_{ij} \forall i \in [N]$), and quantify the "help" that client i needs by $\overline{\Delta \alpha_i}$: the average change of the i -th row of A (i.e. $\alpha_{ij} \forall j \in [N]$). A good indicator of this "help" is the number of local training samples. From the perspective of client j (the helper), larger local training set of client j should be able to make its empirical risk more likely to be helpful to other clients. From the perspective of client i (the helpee), smaller training set of client i should be able to require more help from others. Therefore, ideally, $\overline{\Delta A_j}$ and $\overline{\Delta \alpha_i}$ should be positively and negatively correlated to the number of local training samples, respectively. This is verified as a finding in the resulting A of PGFed in Fig. 4b and Fig. 4c.

B.3. Experiments on Other Datasets

To further evaluate the effectiveness of PGFed with different types of data and different FL settings, we conducted experiments on two other datasets, OrganAMNIST [41] and Office-home [36]. OrganAMNIST is a medical imaging dataset of abdominal CT images in 11 categories. Office-home [36] contains four domains (Art, Clipart, Product, and Real World) of images depicting 65 categories of objects typically found in Office and Home settings.

	25 clients sample 50% Dir(1.0)	50 clients sample 25% Dir(0.3)	100 clients sample 25% Dir(0.3)
Local	90.45±0.19	90.63±0.07	87.14±0.10
FedAvg	99.11±0.03	98.74±0.04	98.47±0.08
APFL	97.49±0.05	97.53±0.06	96.19±0.11
FedRep	95.06±0.16	94.86±0.07	92.47±0.04
LGFedAvg	90.47±0.18	90.99±0.08	87.52±0.22
FedPer	97.89±0.06	97.55±0.08	95.56±0.33
Per-FedAvg	98.40±0.06	96.73±0.03	95.72±0.08
FedRoD	98.61±0.05	98.14±0.09	97.05±0.06
FedBABU	96.49±0.28	94.33±0.13	91.07±0.23
PGFed	99.20±0.04	99.17±0.05	98.94±0.02
PGFedMo	99.21±0.04	99.17±0.07	98.86±0.06

Table 5. Mean and standard deviation over three trials of the mean personalized test accuracy (%) on OrganAMNIST

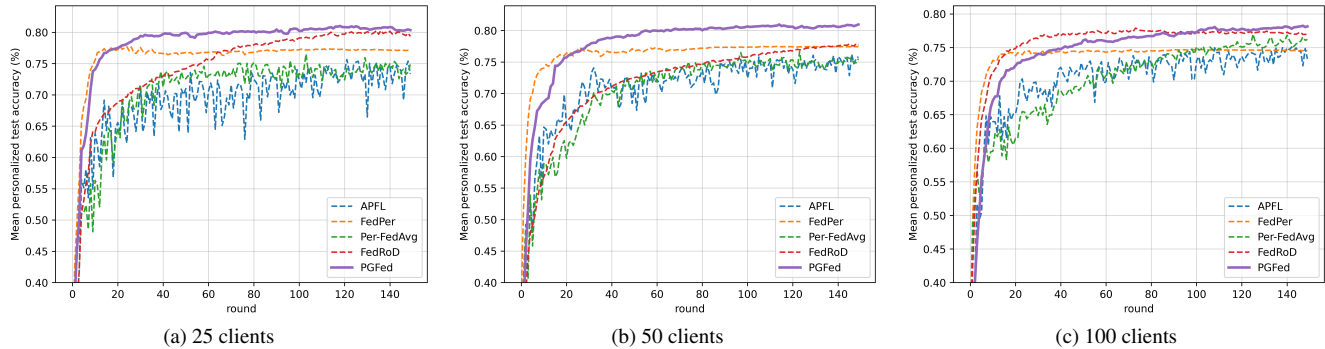


Figure 4. Convergence behavior of the personalized FL approaches with top performance on CIFAR10. While achieving the highest accuracy performance, PGFed is also able to consistently converge faster than several of the baselines that reach high accuracies.




Figure 5. Visualization of the change in \mathbf{A} . Fig. 5a is a heat map of the change in \mathbf{A} . For Fig. 5b and Fig. 5c, the Y-axis of Fig. 5b represents the column average of the change in \mathbf{A} (the average change of weights of client j 's empirical risk on other clients). The Y-axis of Fig. 5c is the row average of the change in \mathbf{A} (the average change of weights of the auxiliary risk on client i). Through the regression line, we verify the positive correlation between $\overline{\Delta \mathbf{A}}_j$ and n_j in Fig. 5b, and the negative correlation between $\overline{\Delta \alpha}_i$ and n_i in Fig. 5c.

For OrganAMNIST, we adopt three settings with different numbers (25, 50, 100) of clients. For the 50- and 100-client settings, we follow the same setting as in the experiments on CIFAR10/CIFAR100, and use the Dirichlet distribution with $\alpha = 0.3$ (Dir(0.3)) and 25% client sample rate for each round. For the 25-client setting, we reduce the heterogeneity in the dataset via Dir(1.0) distribution, and use a higher client sample rate (50%) to simulate a situation more similar to cross-silo FL settings. For Office-home, we adopt a 20-client setting, where each domain contains 5 clients. The non-IIDness in each domain is achieved by Dir(0.3). The mean personalized test accuracies over each domain and over the whole federation are reported. We compare the proposed PGFed and PGFedMo against Local, FedAvg, and the personalized FL baselines that achieved high performance in previous experiments on CIFAR10 and CIFAR100. The results are shown in Tab. 5 and Tab. 6.

For OrganAMNIST, PGFed and PGFedMo achieve the best performance under all three settings. In addition, the proposed algorithms do not have an obvious

drop in the performance from the less heterogeneous 25-client setting to 50-client and 100-client settings. This is not the case for many other personalized FL baselines (FedPer, Per-FedAvg, FedRoD, and FedBABU). Moreover, FedAvg achieves excellent performance on OrganAMNIST, which we believe is due to the similarity of clients' images (see Fig. 6). Since the Dirichlet distribution can only differ the clients in $P(y)$, the label distribution, instead of $P(x|y)$, the distribution of the images given the label, a simple averaging might work unexpectedly well. To see how the methods will perform on different $P(x|y)$ for different clients, we conduct experiments on the Office-home dataset, since for this dataset, even the images within the same class can differ a lot if they belong to clients from different domains. The results on Office-home show that PGFed and PGFedMo achieve the highest mean accuracies over all compared methods. And in each domain, PGFed and PGFedMo consistently outperform most of the compared methods, demonstrating their superiority.

	Art	Clipart	Product	Real World	Mean
Local	17.16 ± 0.85	37.65 ± 0.47	43.83 ± 0.40	24.50 ± 0.21	30.79 ± 0.23
FedAvg	11.68 ± 1.26	41.29 ± 0.85	42.49 ± 1.28	19.14 ± 0.89	28.65 ± 0.49
APFL	19.11 ± 1.55	44.67 ± 0.61	50.40 ± 0.56	25.85 ± 0.88	35.00 ± 0.41
FedRep	20.24 ± 1.45	38.43 ± 1.02	43.70 ± 1.04	24.02 ± 0.81	31.60 ± 0.05
LGFedAvg	17.54 ± 0.45	38.75 ± 0.13	44.59 ± 0.62	25.79 ± 0.61	31.67 ± 0.21
FedPer	17.83 ± 1.07	38.97 ± 0.35	45.87 ± 0.13	25.01 ± 0.52	31.92 ± 0.24
Per-FedAvg	25.90 ± 0.89	42.89 ± 1.77	48.24 ± 1.07	26.00 ± 0.80	35.76 ± 0.21
FedRoD	19.67 ± 1.23	42.44 ± 0.77	44.34 ± 2.07	24.28 ± 1.69	32.68 ± 0.69
FedBABU	18.18 ± 3.54	42.10 ± 2.31	43.51 ± 0.91	26.81 ± 1.86	33.38 ± 0.29
PGFed	<u>22.40 ± 0.26</u>	46.48 ± 1.00	<u>49.86 ± 2.14</u>	26.04 ± 0.80	36.19 ± 0.92
PGFedMo	22.16 ± 0.45	<u>45.88 ± 0.83</u>	49.45 ± 0.19	<u>26.60 ± 0.99</u>	<u>36.02 ± 0.20</u>

Table 6. Mean and standard deviation over three trials of the mean personalized accuracy% of the four domains (5 clients/domain) and the average performance on Office-home dataset. The highest and second-highest accuracies in each column are in **bold** and underlined, respectively.




Figure 6. OrganAMNIST [36] image samples.

C. Comparison in Local Computation Speed

In this section, we study the local computational speed of PGFed and the baselines that achieved top performance in the experiments. We measure the local computational speed by the number of images each method can process per second. We report the local computational speed of the methods in Tab. 7 on CIFAR10 with 50 clients and a batch size of 128, using an NVIDIA Tesla V100 GPU and an Intel(R) Xeon(R) Gold 6248 CPU.

From the results, we can see that PGFed not only reaches high accuracy, but has a relatively high computational speed as well. With a batch size of 128, PGFed reaches a similar local computational speed as FedAvg.

	Images/s	Relative speed	Accuracy
FedAvg	6917.1	100.00%	64.41 ± 0.66
APFL	3389.8	48.99%	77.36 ± 0.18
Per-FedAvg	3464.5	50.09%	76.53 ± 0.31
FedRoD	6682.4	96.61%	79.61 ± 0.22
PGFed	6120.0	88.48%	81.42 ± 0.31
PGFedMo	6032.8	87.22%	81.48 ± 0.32
PGFed-CE*	6175.5	89.28%	81.16 ± 0.56

* A more communication-efficient variation of PGFed, introduced in Appendix D

Table 7. Computational speed (in terms of “images/s”) and accuracy on CIFAR10 with 50 clients

PGFedMo is slightly slower than PGFed due to the momentum update of the auxiliary gradient. However, for some of the compared methods that also achieve high accuracy, their computational speed is compromised by around 50% (compared to FedAvg): APFL needs to train a global model and a local adapter, while Per-FedAvg leverages meta-learning which is a bi-level optimization problem (twice gradient descent on one batch of data). FedRoD trains a global model and a local classifier, which ended up 8.13% faster than PGFed. For PGFed, the extra local computation (over FedAvg) happens at the addition between the gradient from the local empirical risk and the gradient from the auxiliary risk, and also at the update of α_i for each iteration where a dot product of vectorized models is calculated (see Eq. (12) in the main paper).

D. PGFed-CE, a More Communication- and Computation-efficient PGFed

As mentioned in Sec. 6 of the main paper, although PGFed manages to achieve asymptotically the same communication overhead as FedAvg ($O(N)$) while estimating values that, in theory, should take $O(N^2)$ communication cost, since each client is required to download three and upload two models/gradients per round, on average the communication cost is still high (roughly 2.5 times as much as that of FedAvg).

In this section, we provide a more communication-efficient version of PGFed, dubbed PGFed-CE that downloads one less model/gradient from the server. In Sec. 4 of the main paper, we mentioned that $g_\alpha^{(2)}$, a portion of the gradient of the local objective in terms of α_i in PGFed, could be computed on the server instead of the client. This is because

$$g_\alpha^{(2)} = \mu \nabla_{\theta_j} f_j(\theta_j)^T \theta_i, \quad (16)$$

and the server has $\nabla_{\theta_j} f_j(\theta_j)$ from the previous round and the current round global model as the initialization of θ_i , it is possible to treat $g_\alpha^{(2)} = \mu \nabla_{\theta_j} f_j(\theta_j)^T \theta_{glob}$ as a constant computed by the server, where the global model is used as an estimation of the θ_i for the whole round. Since α_i should change adaptively according to the change of θ_i , using the global model as a fixed estimation is not ideal. Nonetheless, this variation saves the communication cost by one less gradient (\bar{g}_{S_i}) from clients' round-beginning download, that *was* needed to adaptively compute the α_i locally, which also slightly saves the local computation. We name such a variation of the PGFed as PGFed-CE. Since each client is now required to download two, instead of three, models/gradients per round, on average the communication cost is reduced from 2.5 to 2 times as much as that of FedAvg. We report the local computational speed of PGFed-CE and its performance on CIFAR10 with 50 clients in Tab. 7. As expected, besides the reduced communication cost, PGFed-CE also simultaneously increases the local computational speed with little drop in the performance.

E. Hyperparameters

Besides the hyperparameter tuning of PGFed reported in Sec. 4.1 in the main paper, we further report the hyperparameter tuning of the compared baselines in this section. The learning rate of all baselines is tuned from $\{0.1, 0.01, 0.001, 0.0001\}$. For FedDyn, we tuned the α in $\{0.1, 0.01, 0.001\}$. For APFL, the α is tuned in $\{0.25, 0.50, 0.75\}$. For Per-FedAvg, the two learning rates for each step are also selected from $\{0.1, 0.01, 0.001, 0.0001\}$, which is the same for FedBABU's FL training and fine-tuning step's learning rates.