

Look Back for More: Harnessing Historical Sequential Updates for Personalized Federated Adapter Tuning

Danni Peng¹, Yuan Wang¹, Huazhu Fu¹, Jinpeng Jiang², Yong Liu¹, Rick Siow Mong Goh¹, Qingsong Wei¹

¹Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A*STAR), Singapore
²EVYD Technology

dannip@ihpc.a-star.edu.sg, wang_yuan@ihpc.a-star.edu.sg, hzfu@ieee.org, jinpeng.jiang@evydtch.com, liuyong@ihpc.a-star.edu.sg, gohsm@ihpc.a-star.edu.sg, wei_qingsong@ihpc.a-star.edu.sg

Abstract

Personalized federated learning (PFL) studies effective model personalization to address the data heterogeneity issue among clients in traditional federated learning (FL). Existing PFL approaches mainly generate personalized models by relying solely on the clients' latest updated models while ignoring their previous updates, which may result in suboptimal personalized model learning. To bridge this gap, we propose a novel framework termed $pFedSeq$, designed for personalizing adapters to fine-tune a foundation model in FL. In $pFedSeq$, the server maintains and trains a sequential learner, which processes a sequence of past adapter updates from clients and generates calibrations for personalized adapters. To effectively capture the cross-client and cross-step relations hidden in previous updates and generate high-performing personalized adapters, $pFedSeq$ adopts the powerful selective state space model (SSM) as the architecture of sequential learner. Through extensive experiments on four public benchmark datasets, we demonstrate the superiority of $pFedSeq$ over state-of-the-art PFL methods.

Introduction

In recent years, federated learning (FL) (McMahan et al. 2017) has attracted growing research interest for enabling privacy-preserving collaborative model training. However, due to data heterogeneity among clients (i.e. data from different clients are non-IID or unbalanced), it is difficult to develop a one-fits-all global model that performs well on all clients' local distributions. To address this, personalized federated learning (PFL) (Smith et al. 2017) has emerged. Unlike traditional FL, which develops a single best model for the collective goal, PFL allows each client to have a unique personalized model tailored specifically to the local objective (Li et al. 2021a; Fallah, Mokhtari, and Ozdaglar 2020). Through the collaborative learning scheme, PFL further enables the personalized models to benefit from knowledge sharing across clients, striking a balance between individualization and generalization for enhanced local performance (Kulkarni, Kulkarni, and Pant 2020; Kairouz et al. 2021).

Recently, as large foundation models (FMs) demonstrate impressive capabilities across various tasks, there has been a rise in research synergizing FL and FMs (Zhuang, Chen,

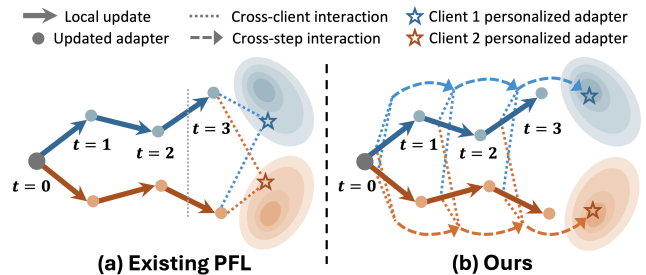


Figure 1: Comparison between (a) existing PFL methods and (b) our approach. Instead of leveraging only the latest updates, our approach accounts for past learning trajectories by modeling cross-client and cross-step relations in previous steps, providing a broader view for identifying the consistent trends for learning more robust personalized adapters.

and Lyu 2023). One of the most popular ways of empowering FL with FMs is through parameter-efficient adapter tuning (e.g., LoRA (Hu et al. 2021)), also known as federated adapter tuning (Li et al. 2024; Woisetschläger et al. 2024). This new FL paradigm allows clients to leverage a powerful, pre-trained FM while only fine-tuning and sharing the lightweight adapters with the server, enabling enhanced performance with minimal on-client computation and communication costs. It also opens up opportunities to tailor PFL methods for fine-tuning FM with adapters. For example, leveraging FL for collaborative fine-tuning while personalizing the adapters to better align the representations of pre-trained FM with the local needs of clients (Yi et al. 2023; Xie et al. 2024; Yang et al. 2024).

However, a common limitation of existing PFL and efforts tailored for FM adapter tuning is that they rely solely on the updates received in the most recent round to develop personalized models, while the potentially valuable information contained in previous updates is either underutilized or completely discarded. This can easily lead to noisy and suboptimal solutions, jeopardizing performance when applied to federated adapter tuning personalization. If we view the development of a personalized adapter for a client as an optimization process, focusing only on the latest updates is akin to performing simple gradient descents. By incorporating information from previous updates—much like using mo-

mentum in optimization—the learning process accounts for a broader range of past trajectories and uncovers the consistent trends, leading to more robust and reliable personalized solutions at convergence. Hence, in this work, we are motivated to design a framework that accounts for past learning trajectories to develop enhanced personalized adapters. To illustrate, Figure 1 shows the adapter update trajectories of two clients over three rounds of local tuning and communication. As depicted in Figure 1a, existing PFL methods typically rely only on the interactions among the latest updates to produce personalized adapters, which can lead to less robust and undeterministic solutions by only referencing the current states. In contrast, our approach aims to model the interactions among clients at each past update step, as well as the dependencies across various steps, as shown in Figure 1b, which provides a broader perspective for identifying the consistent update patterns, enabling more robust learning and development of more superior personalized adapters.

Specifically, we propose a novel PFL framework termed $p\text{FedSeq}$ for improving personalized federated adapter tuning by exploiting knowledge from clients’ past sequential updates. This is achieved through two main processes at the server during each communication round: (1) standard model aggregation using a traditional FL algorithm to obtain a global adapter (e.g., FedAvg (McMahan et al. 2017)), and (2) leveraging a *Sequential Learner* to process the sequence of clients’ adapter updates collected at the server and output personalized calibrations to adjust the global adapter for clients’ individual needs. To enable flexible modeling of the cross-client and cross-step relations in the sequence, we employ a learnable hypernetwork for the sequential learner, which is jointly trained at the server by optimizing the personalized adapters toward clients’ local objectives. By using the received adapter updates as a proxy for gradients on local losses, the training of sequential learner can be efficiently performed at server without accessing clients’ local data.

Selecting a suitable architecture for the sequential learner is a non-trivial task. To facilitate effective capture of cross-client and cross-step relations, we propose to employ the selective state space model (SSM) (Gu and Dao 2023) as our sequential learner. Selective SSM is a recurrence-based model introduced recently for efficient sequence modeling. It enjoys both effective performance with time-dependent selectivity and efficient computations with linear scaling. To capitalize on its design features, we form the step-wise inputs by concatenating clients’ updates from the same round and processing them in a recurrence mode. With that, the cross-client interactions at different steps can be captured in the time-dependent module parameters, and the sequential processing allows the cross-step dependencies to be captured in the consolidated hidden states.

To summarize, our contributions are as follows:

- We propose a novel PFL framework, $p\text{FedSeq}$, for personalizing federated adapter tuning for clients. By leveraging knowledge from the previous updates, $p\text{FedSeq}$ generates enhanced personalized adapters which better tailor large FMs’ representations to clients’ local needs.
- We collaboratively train a sequential learner at the server

to capture useful cross-client and cross-step relations in the sequential updates, achieving effective personalized adapter generation by adopting Selective SSM as the learner architecture.

- We evaluate our $p\text{FedSeq}$ rigorously on four large-scale benchmark datasets (i.e., CIFAR-100, Tiny-ImageNet, DomainNet, and Omniglot), and show that our $p\text{FedSeq}$ outperforms ten state-of-the-art PFL methods by up to 5.39%.

Related Work

Traditional FL algorithms (McMahan et al. 2017; Li et al. 2020; Karimireddy et al. 2020) following the one-model-fits-all paradigm often suffer from degraded performance on clients’ local data in the face of severe data heterogeneity. Recently, personalized FL (PFL) has attracted much attention, which develops customized models to accommodate the diverse needs of clients (Mansour et al. 2020). Generally, efforts in PFL fall under four categories:

(1) Meta-learning-based methods. By drawing an analogy between FL and meta-learning (Jiang et al. 2019), this approach jointly develops a global model from which the personalized models can be effectively fine-tuned with just a few local steps. Per-FedAvg (Fallah, Mokhtari, and Ozdaglar 2020) adopts the spirit of MAML (Finn, Abbeel, and Levine 2017) and jointly learns a global initialization by optimizing for one-step gradient updates. $p\text{FedMe}$ (T Dinh, Tran, and Nguyen 2020) further allows multiple updates in the inner loop and optimizes a Moreau envelope objective.

(2) Personalized-aggregation-based methods. This line of methods produces personalized models by learning client-specific aggregation weights to combine models from other clients. FedFomo (Huang et al. 2021) and FedAMP (Zhang et al. 2021) leverage a rule-based approach to compute weights based on the pair-wise distance between clients’ models. APPLE (Luo and Wu 2022) and FedALA (Zhang et al. 2023) adopt a learning-based approach to optimize the personalized weights directly on clients’ local objectives. FedDPA (Yang et al. 2024) was introduced recently, focusing on addressing test-time distribution shifts by combining global and personalized adapters through dynamic weighting for federated adapter tuning.

(3) Personalized-network-based methods. This approach develops a personalized model independently for each client while drawing on the general knowledge through various forms of global sharing. FedRep (Collins et al. 2021) and FedBN (Li et al. 2021b) locally update parts of the network that are sensitive to data distributions (e.g., the classification head or the batch-norm layers), while sharing the rest to leverage the common representations. Ditto (Li et al. 2021a) trains full personalized models locally, using a proximal term to regularize distance from the global model. PerAda (Xie et al. 2024) was recently introduced for federated adapter tuning. It leverages proximal regularization similar to Ditto, while further incorporating knowledge distillation to facilitate information sharing.

pFedLoRA (Yi et al. 2023) encourages efficient sharing for model-heterogeneous PFL through lightweight adapters.

(4) Hypernetwork-based methods. As opposed to solely relying on model aggregation for information sharing, this approach outlines a new way of federation by jointly training a hypernetwork using feedback returned from clients (e.g., the model updates). The hypernetwork is trained to directly generate personalized model parameters based on some inputs. Since it is only trained and maintained at the server, a high-capacity, complex network can be adopted for enhanced diversity of the models generated without concern about the communication costs. pFedHN (Shamsian et al. 2021) is a pioneering work that leverages this approach to generate personalized models for clients based on the learnable client descriptor vectors. L2C (Li et al. 2022) and pFedLA (Ma et al. 2022) utilize hypernetwork to learn personalized aggregation weights. pFedPG (Yang, Wang, and Wang 2023) focuses on federated prompt learning and trains a hypernetwork for personalized prompt generation. PeFLL (Scott, Zakerinia, and Lampert 2024) and FedL2P (Lee et al. 2023) condition on clients’ local statistics and train a hypernetwork to output personalized models or update strategies.

Our work, by collaboratively training a sequential learner, falls under the hypernetwork-based methods. Different from the existing works, we propose to generate personalized adapters by leveraging clients’ previous updates. In other related fields concerning multi-task/distribution learning similar to FL, past gradient updates are commonly used to enhance task representations (Zenke, Poole, and Ganguli 2017; Flennerhag et al. 2018; Peng and Pan 2023) or to derive relations among different parties (Yu et al. 2020; Mansilla et al. 2021). In the realm of FL, (Ji et al. 2019) utilizes previous clients’ updates to produce global model via an RNN-based aggregator. However, the potential of leveraging previous updates for PFL still remains unexplored. Our work fills this gap by developing a method to extract useful cross-client and cross-step relations from past updates, producing personalized adapters that better suit clients’ local specifics.

Method

Preliminaries

Problem Setup. In a typical FL setup involving N clients and a central server, each client $i \in [N]$ has its own private data \mathcal{D}_i . Traditional FL (e.g., FedAvg (McMahan et al. 2017)) seeks to learn a global model ϕ_g that performs well across all clients: $\min_{\phi_g} \sum_{i=1}^N \mathcal{L}(\mathcal{D}_i; \phi_g)$, where \mathcal{L} is an arbitrary loss function. However, this one-model-fits-all scheme may fail when data heterogeneity is severe among clients. PFL addresses this by relaxing the single-model constraint and learning N personalized models $\{\phi_i\}_{i=1}^N$, each tailored specifically to a client, while also benefiting from the federation by drawing knowledge from global model ϕ_g :

$$\min_{\{\phi_i\}_{i=1}^N, \phi_g} \sum_{i=1}^N \mathcal{L}(\mathcal{D}_i; \phi_i, \phi_g).$$

Personalized Federated Adapter Tuning. Adopting adapter tuning in PFL provides a computation- and communication-efficient solution for clients to harness the

power of large FMs. Let $F^*(\cdot)$ denote a fixed, pre-trained foundation model backbone (e.g., ViT (Dosovitskiy et al. 2020)), which remains locally at clients. A chosen adapter (e.g., LoRA), denoted by $a_\theta(\cdot)$ parameterized by θ , as well as the classification head $h_\omega(\cdot)$ parameterized by ω , are tuned at the client to adapt the fixed backbone to local distributions. Since only the adapter and the head are updated and shared with the server, without loss of generality, we define the global and personalized models as $\phi_g = (\theta_g, \omega_g)$ and $\phi_i = (\theta_i, \omega_i)$. The individual loss of client i is computed by:

$$\mathcal{L}(\mathcal{D}_i; \phi_i, \phi_g) := \sum_{(x,y) \in \mathcal{D}_i} l(h_{\omega_i, \omega_g}(F^*(x) + a_{\theta_i, \theta_g}(x)), y).$$

pFedSeq Framework

In this section, we present our pFedSeq framework, designed for personalizing federated adapter tuning. Following (Collins et al. 2021; Yang, Wang, and Wang 2023), we update the classification head locally without sharing with the server to better preserve the client-specific knowledge, and apply pFedSeq to generate the personalized adapters.

An overview of pFedSeq is shown in Figure 2. At each communication round, pFedSeq performs *local adapter tuning* at the clients and sends the adapter updates to the server. The server then conducts *personalized adapter generation* with the sequential learner by processing the sequence of past updates collected from clients. Meanwhile, the server also performs *sequential learner optimization* using the new updates received from the clients as feedback for training. We now introduce each process in detail.

Local Adapter Tuning at Client. Suppose we are at the t -th local training round of client i . Upon receiving the personalized adapter θ_i^{t-1} generated by the server in the previous round, we update θ_i^{t-1} and ω_i^{t-1} jointly on local data \mathcal{D}_i for several local epochs to obtain $\tilde{\theta}_i^t$ and ω_i^t . Note that ω_i^{t-1} is not generated by the server and is restored from the previous local update at client i . Also, to be differentiated from the personalized adapter θ_i^t generated by the server, we use $\tilde{\theta}_i^t$ to denote the adapter updated locally. We then compute the adapter update $\Delta_i^t := \tilde{\theta}_i^t - \theta_i^{t-1}$ and send it to the server.

Personalized Adapter Generation at Server. To generate personalized adapters, two processes are carried out at the server: (1) aggregating clients’ updated adapters to produce a global adapter, and (2) using the sequential learner to process clients’ past updates and output personalized calibrations for tailoring the global adapter to each client.

Specifically, after the t -th communication round, server receives adapter updates $\{\Delta_i^t\}_{i=1}^N$ from N clients. To compute the globally aggregated adapter, we need to obtain the clients’ updated adapters $\{\tilde{\theta}_i^t\}_{i=1}^N$. To avoid doubling the communication costs, we can compute $\{\tilde{\theta}_i^t\}_{i=1}^N$ directly using the adapter updates received from the clients and the personalized adapters $\{\theta_i^{t-1}\}_{i=1}^N$ generated at the server in the previous round, i.e., $\tilde{\theta}_i^t = \theta_i^{t-1} + \Delta_i^t, \forall i \in [N]$. By using the classic FedAvg (McMahan et al. 2017) for aggregation, we obtain the global adapter $\tilde{\theta}_g^t$ at the t -th round:

$$\tilde{\theta}_g^t = \sum_{i=1}^N \frac{|\mathcal{D}_i|}{\sum_{i'=1}^N |\mathcal{D}_{i'}|} \tilde{\theta}_i^t. \quad (1)$$

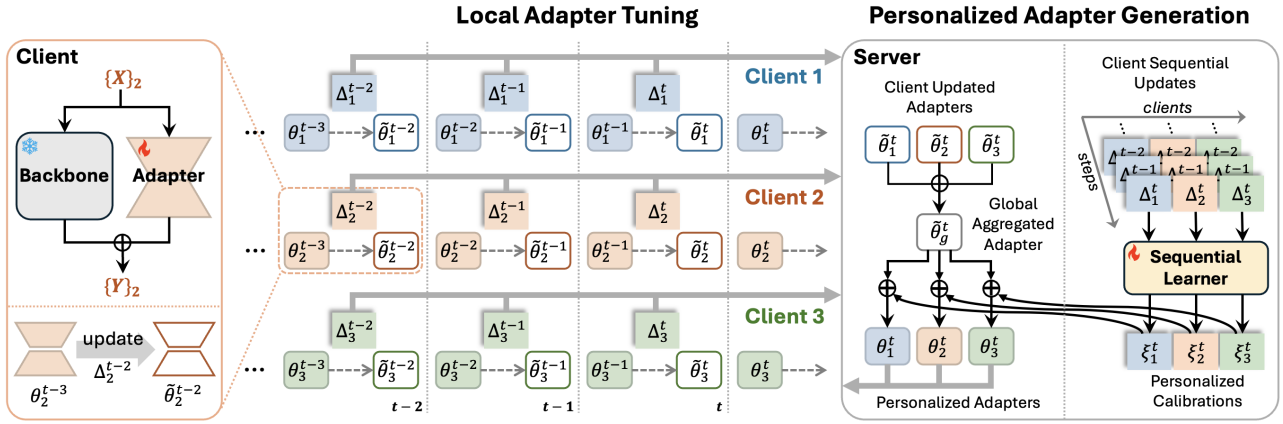


Figure 2: An overview of pFedSeq framework. At each communication round t , clients perform local adapter tuning and send the adapter updates $\{\Delta_i^t\}_{i=1}^N$ to server. At the server, the updated adapters are aggregated to form a global adapter $\tilde{\theta}_g^t$. Meanwhile, the sequential learner processes the sequence of updates collected at the server and generates personalized calibrations, which are then applied to the global adapter to produce the personalized adapters $\{\theta_i^t\}_{i=1}^N$ and send to the clients.

Simultaneously, we utilize a sequential learner, which is a hypernetwork, to generate personalized calibrations from the sequential updates. Specifically, at the t -th round, we construct the input to sequential learner by first stacking $\{\Delta_i^j\}_{i=1}^N$ across N clients for each $j \in [t]$ collected at the server, forming $\Delta^j = [\Delta_1^j, \dots, \Delta_N^j] \in \mathbb{R}^{D \times N}$, where D is the dimensionality of the adapter's parameters. We then concatenate $\{\Delta^j\}_{j=1}^t$ across t steps, forming the sequence input matrix $\Delta^{1:t} = [\Delta^1, \dots, \Delta^t] \in \mathbb{R}^{D \times N \times t}$. The sequential learner, denoted by $\text{SeqLearner}(\cdot; \psi)$ parameterized by ψ , outputs the personalized calibrations $\xi^t = [\xi_1^t, \dots, \xi_N^t] \in \mathbb{R}^{D \times N}$ for the N clients by taking in the sequence input $\Delta^{1:t}$:

$$\xi^t = \text{SeqLearner}(\Delta^{1:t}; \psi). \quad (2)$$

Note that by treating D as the batch dimension, we employ one $\text{SeqLearner}(\cdot; \psi)$ to capture the cross-client and cross-step relations for all parameters of an adapter. Hence, the size of $\text{SeqLearner}(\cdot; \psi)$ is independent of the adapter size D and only depends on the number of clients N and the sequence length t . For better expressivity, we assign one $\text{SeqLearner}(\cdot; \psi)$ to learn the parameters of the adapter attached to each layer of the backbone (e.g., ViT-B/16 contains 12 layers (Dosovitskiy et al. 2020)).

Finally, we obtain the personalized adapters $\{\theta_i^t\}_{i=1}^N$ for the t -th round by adding $\{\xi_i^t\}_{i=1}^N$ to the global adapter $\tilde{\theta}_g^t$:

$$\theta_i^t = \tilde{\theta}_g^t + \xi_i^t, \quad \forall i \in [N]. \quad (3)$$

The personalized adapters $\{\theta_i^t\}_{i=1}^N$ are sent to the clients for the next round of local training.

Sequential Learner Optimization at Server. Since the objective of the sequential learner is to generate effective personalized calibrations that perform well on clients' local data, we optimize the parameters ψ of $\text{SeqLearner}(\cdot; \psi)$ by:

$$\min_{\psi} \sum_{i=1}^N \mathcal{L}(\mathcal{D}_i; \theta_i), \quad (4)$$

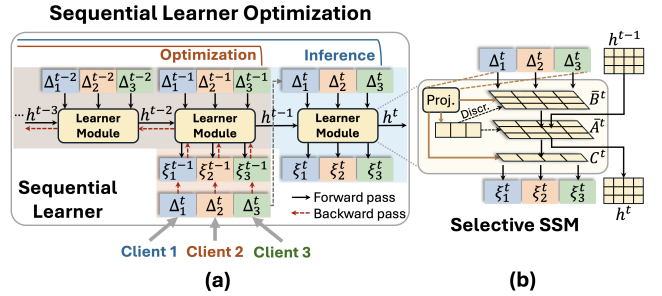


Figure 3: (a) Optimization and inference processes of sequential learner at the t -th round. (b) An instantiation of sequential learner using Selective SSM.

where $\theta_i = \tilde{\theta}_g + \xi_i = \tilde{\theta}_g + (\text{SeqLearner}(\Delta; \psi))_{:,i}$. Using chain rule, the gradient update for ψ from each client i is given by $\nabla_{\psi} \mathcal{L}(\mathcal{D}_i; \theta_i) = (\nabla_{\psi} \xi_i)^{\top} \nabla_{\theta_i} \mathcal{L}(\mathcal{D}_i; \theta_i)$ (note that $\nabla_{\xi_i} \theta_i = \mathbb{I}$). Following (Shamsian et al. 2021; Scott, Zakerinia, and Lampert 2024), we approximate $\nabla_{\theta_i} \mathcal{L}(\mathcal{D}_i; \theta_i)$ using the adapter update Δ_i received from client i . This is equivalent to replacing a single gradient step on θ_i with multiple gradient update steps, which has been shown to achieve better and more stable convergence (Shamsian et al. 2021).

Specifically, at the t -th round, we receive adapter updates $\{\Delta_i^t\}_{i=1}^N$ from clients. Recall that these are the updates from the previous round's personalized adapters $\{\theta_i^{t-1}\}_{i=1}^N$, i.e., Δ_i^t is the proxy for gradient evaluated at θ_i^{t-1} , indicating the optimization direction at θ_i^{t-1} . Hence, we use $\{\Delta_i^t\}_{i=1}^N$ as the signal to adjust $\{\theta_i^{t-1}\}_{i=1}^N$, which constitutes $\{\xi_i^{t-1}\}_{i=1}^N$ generated by inputting $\Delta^{1:t-1}$ (i.e., sequence until $t-1$) to $\text{SeqLearner}(\cdot; \psi)$. Formally, we compute the update $(\Delta\psi)^t$ for ψ at the t -th round by:

$$(\Delta\psi)^t = \sum_{i=1}^N (\nabla_{\psi} \xi_i^{t-1})^{\top} \Delta_i^t, \quad (5)$$

where $\xi_i^{t-1} = (\text{SeqLearner}(\Delta^{1:t-1}; \psi))_{:,i}$.

After updating $\text{SeqLearner}(\cdot; \psi)$, we perform inference to generate $\{\xi_i^t\}_{i=1}^N$ by inputting $\Delta^{1:t}$ (including the latest update Δ^t) to $\text{SeqLearner}(\cdot; \psi)$, as described in (2). The optimization and inference processes of sequential learner at the t -th round are illustrated in Figure 3a.

Remarks. To avoid infinitely growing sequence length as the number of update rounds increases, we cap the sequence input at a maximum length L , i.e., (2) becomes:

$$\xi^t = \begin{cases} \text{SeqLearner}(\Delta^{1:t}; \psi) & \text{for } t \leq L \\ \text{SeqLearner}(\Delta^{t-L+1:t}; \psi) & \text{for } t > L \end{cases} \quad (6)$$

Also, to ensure that $\text{SeqLearner}(\cdot; \psi)$ generates reliable personalized adapters for the next round of local update for stable convergence, we set a warm-up period W to sufficiently train $\text{SeqLearner}(\cdot; \psi)$ before putting it into use. That is, for the first W rounds, the server only updates $\text{SeqLearner}(\cdot; \psi)$ without using it to generate the personalized calibrations, and only the global adapter $\hat{\theta}_g^t$ is sent to clients during the warm-up period. Algorithm 1 in Appendix A summarizes the workflow.

Sequential Learner using Selective SSM

In this section, we introduce an instantiation of the sequential learner using Selective SSM as the learner architecture. Selective SSM is recently introduced for efficient sequence modeling (Gu and Dao 2023). In terms of modeling cross-client and cross-step relations, the selection mechanism of Selective SSM allows cross-client interactions at different steps to be captured in the input-dependent parameters. At the same time, the recursive processing effectively abstracts the cross-step dependencies in the internal hidden states.

Given a sequence input $\Delta^{t-L+1:t} \in \mathbb{R}^{D \times N \times L}$ of length L , Selective SSM generates output $\xi^j \in \mathbb{R}^{D \times N}$ at each step $j \in [t-L+1 : t]$ by taking in the previous step hidden state $h^{j-1} \in \mathbb{R}^{D \times N \times M}$ and the current step input $\Delta^j \in \mathbb{R}^{D \times N}$, where M is an expanded latent dimension. The step-wise modular operation is formulated as follows:

$$h^j = \bar{\mathbf{A}}^j h^{j-1} + \bar{\mathbf{B}}^j \Delta^j, \quad \xi^j = \mathbf{C}^j h^j, \quad (7)$$

where $\bar{\mathbf{A}}^j, \bar{\mathbf{B}}^j \in \mathbb{R}^{D \times N \times M}$ are discretized parameters, and $\mathbf{C}^j \in \mathbb{R}^{D \times M}$ is a projection matrix, all obtained by conditioning on the current step input Δ^j (a concatenation of N clients' updates), i.e., $\bar{\mathbf{A}}^j = s_{\bar{\mathbf{A}}}(\Delta^j)$, $\bar{\mathbf{B}}^j = s_{\bar{\mathbf{B}}}(\Delta^j)$, $\mathbf{C}^j = s_{\mathbf{C}}(\Delta^j)$ (detailed formulations in (Gu and Dao 2023)). Note that we input a sequence $\Delta^{t-L+1:t}$ into Selective SSM, and optimize or perform inference only on the final-step output ξ^t . Figure 3b illustrates a step module of Selective SSM.

Following (Gu and Dao 2023), we incorporate a 1D convolution and a residual connection before and after the Selective SSM, forming a Mamba block. A detailed description of our architecture is included in Appendix B.

Experiments

Experimental Setup

Datasets and Heterogeneity Scenarios. We evaluate our pFedSeq on four benchmark datasets covering three different data heterogeneity scenarios. For *label-skew* scenario,

we use **CIFAR-100** (Krizhevsky et al. 2009) which consists of 60,000 images from 100 classes, and **Tiny-ImageNet** (Chrabaszcz, Loshchilov, and Hutter 2017) which consists of 110,000 images from 200 classes. For both datasets, we simulate label-skew heterogeneity by distributing data in each class over 10 clients with a Dirichlet distribution $Dir(0.1)$, following (Yang, Wang, and Wang 2023; Zhang et al. 2023). For *feature-skew* scenario, we consider **DomainNet**, which involves 600,000 images from 345 classes across 6 domains (i.e., Clipart, Infograph, Painting, Quickdraw, Real, and Sketch). Following (Li et al. 2021b), we use the top ten most frequent classes for experiments, simulating feature skew across clients by treating each domain as a client (i.e., $N = 6$). Furthermore, we consider a *real-world* heterogeneous scenario where data are collected from actual clients. We adopt **Omniglot**, which contains images of 1,623 characters from 50 alphabets, handwritten by 20 different individuals. We treat each individual as a client (i.e., $N = 20$) and predict the alphabet to which a character belongs.

Baselines. We compare the performance of pFedSeq against 12 baselines, including 2 traditional methods and 10 state-of-the-art PFL methods. For traditional baselines, we consider **Local** and classic **FedAvg**, where the former performs local adapter tuning only without sharing information with other clients, while the latter aggregates a global adapter and a global head to share with all clients. For PFL baselines, we include meta-learning-based methods **Per-FedAvg** (Fallah, Mokhtari, and Ozdaglar 2020) and **pFedMe** (T Dinh, Tran, and Nguyen 2020); personalized-aggregation-based methods **APPLE** (Luo and Wu 2022) and **FedALA** (Zhang et al. 2023); personalized-network-based methods **FedRep** (Collins et al. 2021), **Ditto** (Li et al. 2021a), and **PerAda** (Xie et al. 2024); and hypernetwork-based methods **pFedHN** (Shamsian et al. 2021), **pFedLA** (Ma et al. 2022), and **PeFLL** (Scott, Zakerinia, and Lampert 2024). All methods are evaluated on clients' local test sets, and the final result is computed by averaging over all clients.

Implementation Details. For fair comparisons, we adapt all compared methods to the same federated adapter tuning setup, where we adopt ViT-B/16 (Dosovitskiy et al. 2020) pre-trained on ImageNet21k (Deng et al. 2009) as the fixed backbone F^* and fine-tune LoRA adapter (Hu et al. 2021). We implement all methods by tuning and sharing only the LoRA and the classification head using the respective PFL algorithms. For all datasets, the number of communication rounds T is set to 80. At each round, the clients perform adapter tuning for 1 local epoch using SGD optimizer with a batch size of 32. The local learning rate is set to 0.05 for Omniglot and 0.005 for other datasets. For our pFedSeq, we adopt a 2-layer Mamba as our sequential learner and set the expanded state dimension M to 16. We train our sequential learner using Adam optimizer with learning rate 0.001, similarly for other hypernetwork-based methods. For all datasets, we set the number of warm-up rounds W to 10 and tune the maximum sequence length L in $\{5, 10, 15, 20, 25, 30\}$. All experiments are conducted on NVIDIA A100 GPUs with 40GB memory. We repeat each experiment with 3 seeds and report the mean and standard deviation. More

Method	Label-Skew		Feature-Skew	Real-World
	CIFAR-100	Tiny-ImageNet	DomainNet	Omniglot
Local	92.94±0.02	92.52±0.11	78.67±0.15	38.36±0.24
FedAvg	87.36±0.12	87.98±0.19	78.19±0.19	37.87±0.84
<i>Meta-Learning-Based</i>				
Per-FedAvg	93.68±0.08	93.28±0.03	80.62±0.67	40.73±0.40
pFedMe	93.35±0.10	93.62±0.04	81.71±0.28	41.90±0.44
<i>Personalized-Aggregation-Based</i>				
APPLE	93.38±0.08	93.23±0.01	80.34±0.81	40.35±0.72
FedALA	93.73±0.02	93.30±0.02	81.17±0.76	39.86±0.46
<i>Personalized-Network-Based</i>				
FedRep	94.06±0.11	93.01±0.18	81.31±0.48	38.04±0.10
Ditto	93.87±0.05	93.46±0.02	80.74±0.90	41.18±0.46
PerAda	93.91±0.07	93.43±0.02	80.93±0.85	41.72±0.53
<i>Hypernetwork-Based</i>				
pFedHN	94.46±0.37	93.51±0.18	81.98±0.40	42.48±0.19
pFedLA	93.42±0.21	93.37±0.11	80.26±0.20	41.27±0.55
PeFLL	94.41±0.05	93.70±0.26	82.42±0.17	42.98±0.18
pFedSeq	95.30±0.09	94.30±0.08	84.63±0.24	45.25±0.16

Table 1: Performance comparison on four datasets. For fair comparisons, all methods are adapted to the same federated adapter tuning setup (i.e., fine-tuning fixed ViT with LoRA).

implementation details can be found in Appendix C.

Baseline Comparison

Table 1 compares the performance of pFedSeq against baselines applied to the same federated adapter tuning setup. On all the four datasets, our pFedSeq achieves the highest performance, with significant margins of {0.84%, 0.6%, 2.21%, 2.27%} over the second-best performer (i.e., pFedHN or PeFLL). Notably, we observe greater gains for DomainNet and Omniglot characterized by distinct domain discrepancies across clients. This demonstrates the efficacy of our method in personalizing adapters to better adjust the backbone representations for different styles or feature distributions. Throughout Table 1, we observe that all PFL methods outperform Local and FedAvg, indicating their effectiveness in balancing global and local knowledge. As compared to meta-learning and personalized-network methods, which drive the personalized adapters too closely to the global model (through common initialization or regularization), our pFedSeq directly produces diverse personalized calibrations that better adjust the global adapter to clients’ local specifics, leading to better results. Unlike aggregation-based methods, which linearly combine client-specific adapters, pFedSeq leverages the non-linearity of a hypernetwork to model the complex client relations and achieves better knowledge transfer. A closer look at Table 1 reveals that hypernetwork-based methods, such as pFedHN, PeFLL, and our pFedSeq, outshine other types of methods, demonstrating the advantages of using hypernetworks to directly produce small-sized adapters. Our pFedSeq, by using the sequential learner to capture the cross-client and cross-step relations, further outperforms the existing




Figure 4: Learning curves of pFedSeq and compared baselines for CIFAR-100 and Tiny-ImageNet.

hypernetwork-based methods. More discussions on baseline comparisons are included in Appendix D.

In Figure 4, we show the learning curves of pFedSeq and compared baselines over 80 communication rounds for CIFAR-100 and Tiny-ImageNet (the plots for DomainNet and Omniglot are included in Appendix D). For clearer visualization, we only present the curves of representative baselines under each PFL category. From the plots, we can see that our pFedSeq (blue line) begins to outperform all the baselines at around the 25-th round. Though our pFedSeq shows similar performance to FedRep (green line) during the warm-up phase (i.e., the first 10 rounds), where only the global adapter is sent out for clients’ local evaluation, we observe rapid improvement once pFedSeq begins to use the personalized adapters generated by the stably trained sequential learner, quickly surpassing all baselines in the next 15 rounds. The accelerated learning demonstrates the advantages of leveraging previous update steps for deriving more robust and superior personalized adapters, which enables a positive feedback loop between enhanced local updates and better generated personalized adapters, facilitating faster learning and improved performance at convergence.

Analysis of pFedSeq

In this section, we first conduct experiments to verify the effectiveness of the various design components of pFedSeq. Then, we provide detailed analysis of the impact of the maximum sequence length L on the performance of pFedSeq.

Effectiveness of Key Components. We examine how the three key components of pFedSeq (i.e., global aggregation, cross-step modeling, cross-client modeling) contribute to its overall performance by introducing three variants: variant A removes global aggregation from pFedSeq and uses the sequential learner to generate personalized adapters directly instead of personalized calibrations; variant B eliminates cross-step modeling by using only the latest updates from all clients to generate personalized calibrations, equivalent to setting $L = 1$; variant C eliminates cross-client modeling by generating the personalized calibration for a client using only that client’s sequential updates.

As shown in Table 2, we see that removing global aggregation results in the largest performance drops: 0.89% for CIFAR-100 and 1.96% for DomainNet, signifying the importance of directly leveraging the knowledge sharing through the global adapter. Also, we note that even with-

Variant	global aggregation	cross-step modeling	cross-client modeling	CIFAR-100	DomainNet
A	✗	✓	✓	94.41±0.11	82.67±0.26
B	✓	✗	✓	94.66±0.08	83.10±0.36
C	✓	✓	✗	94.40±0.12	83.85±0.23
pFedSeq	✓	✓	✓	95.30±0.09	84.63±0.24

Table 2: Performance of variants with individual key components removed from pFedSeq.

Variant	Learner Architecture	CIFAR-100	DomainNet
D	MLP	94.32±0.17	81.25±0.56
E	LSTM	94.69±0.23	82.77±0.36
pFedSeq	Selective SSM	95.30±0.09	84.63±0.24

Table 3: Performance of variants using different architectures for the sequential learner.

out global aggregation, variant A performs comparably to the strongest baselines (e.g., pFedHN and PeFLL), which shows the effectiveness of the cross-step and cross-client modeling for directly generating the personalized adapters. Next, we see that removing either cross-step or cross-client modeling leads to significant performance drops on both datasets: 0.64% and 0.9% for CIFAR-100, and 1.53% and 0.78% for DomainNet, indicating the importance of taking into account both and exploiting their coupled effects in our design. In addition, we observe that both variants B and C surpass the strongest baselines on the two datasets, which shows that even when considered individually, either cross-step or cross-client modeling is effective for generating enhanced personalized adapters.

Effectiveness of Learner Architecture. To verify our choice of using Selective SSM as the sequential learner, we further introduce two variants using different architectures. First, variant D employs an MLP-based network similar to (Shamsian et al. 2021) for the sequential learner, where the sequence of inputs are concatenated along a single dimension (i.e., $\Delta^{t-L+1:t} \in \mathbb{R}^{D \times (N \times L)}$) and passed into the MLP network. Note that this architecture can only process fixed-length sequences. As shown in Table 3, using an MLP learner leads to performance drops of 0.98% for CIFAR-100 and 3.38% for DomainNet. The decline is mainly attributed to the less effective structure of MLP for modeling sequence inputs, leading to instability of training (as can be observed in Figure 9 in Appendix D). As for the second architecture, variant E utilizes the classical LSTM (Hochreiter and Schmidhuber 1997) for the learner. Similar to Selective SSM, LSTM is capable of modeling variable-length sequences. However, it reuses the same cell for all the steps without selectivity, making it less capable of discerning valuable context information from different steps. Table 3 shows that using an LSTM learner leads to performance declines of 0.61% on CIFAR-100 and 1.86% on Domain-




Figure 5: Performance of pFedSeq by varying the maximum sequence length L on Omniglot.

Net. The less satisfactory performance of variants D and E on two datasets demonstrates the superiority of using Selective SSM to model the cross-client and cross-step dependencies for more effective personalization.

Impact of Maximum Sequence Length L . We investigate the impact of L by varying it in a more fine-grained range $\{1, 3, 5, 7, 9, 11, 13, 15\}$ on Omniglot dataset. The results are plotted in Figure 5, where the dotted line indicates the performance of the strongest baseline PeFLL as a reference. From the plot, we can clearly observe an increasing trend in the performance of pFedSeq as the maximum sequence length L increases. As compared to modeling only on the latest updates (i.e., $L = 1$), modeling on a longer sequence of previous updates (i.e., $L = 15$) results in 1.17% increase in the performance of pFedSeq. This further confirms the effectiveness of our approach in modeling previous sequential updates. Moreover, the performance appears to plateau at a certain level, with further increases in L yielding minimal improvements. This may be because earlier updates that are too distant from the present offer less relevant information, making them less useful for generating personalized adapters for the next update round. To achieve the best trade-off between computational cost and performance, we tune L to find the elbow point where the performance starts to plateau. Overall, our pFedSeq outperforms PeFLL even with $L = 1$, signifying the effectiveness of our choice of architecture in generating personalized calibrations, and our global aggregation process for explicitly leveraging the global knowledge. A plot of the learning curves for different values of L is included in Appendix D.

Conclusion

In this paper, we propose a novel pFedSeq framework for personalizing federated adapter tuning by exploiting knowledge from clients’ previous updates. Our pFedSeq introduces a sequential learner jointly trained across all clients at the server to capture the cross-client and cross-step relations from the sequential updates, and output effective personalized adapters. For the learner architecture, we employ the powerful Selective SSM to leverage its sequence modeling capabilities. Extensive experiments on four benchmark datasets demonstrate the superiority of pFedSeq over ten state-of-the-art PFL methods and verify the effectiveness of its various components through rigorous studies.

Acknowledgments

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Project (IAF-ICP) (Award No: I2301E0020), administered by A*STAR.

References

- Chrabaszcz, P.; Loshchilov, I.; and Hutter, F. 2017. A down-sampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*.
- Collins, L.; Hassani, H.; Mokhtari, A.; and Shakkottai, S. 2021. Exploiting shared representations for personalized federated learning. In *International conference on machine learning*, 2089–2099. PMLR.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Fallah, A.; Mokhtari, A.; and Ozdaglar, A. 2020. Personalized federated learning with theoretical guarantees: a model-agnostic meta-learning approach. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 3557–3568.
- Finn, C.; Abbeel, P.; and Levine, S. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135. PMLR.
- Flennerhag, S.; Moreno, P. G.; Lawrence, N. D.; and Damianou, A. 2018. Transferring Knowledge across Learning Processes. In *International Conference on Learning Representations*.
- Gu, A.; and Dao, T. 2023. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hu, E. J.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.
- Huang, Y.; Chu, L.; Zhou, Z.; Wang, L.; Liu, J.; Pei, J.; and Zhang, Y. 2021. Personalized cross-silo federated learning on non-iid data. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 7865–7873.
- Ji, J.; Chen, X.; Wang, Q.; Yu, L.; and Li, P. 2019. Learning to learn gradient aggregation by gradient descent. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 2614–2620.
- Jiang, Y.; Konečný, J.; Rush, K.; and Kannan, S. 2019. Improving federated learning personalization via model agnostic meta learning. *arXiv preprint arXiv:1909.12488*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z.; Cormode, G.; Cummings, R.; et al. 2021. Advances and open problems in federated learning. *Foundations and trends® in machine learning*, 14(1–2): 1–210.
- Karimireddy, S. P.; Kale, S.; Mohri, M.; Reddi, S.; Stich, S.; and Suresh, A. T. 2020. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, 5132–5143. PMLR.
- Krizhevsky, A.; et al. 2009. Learning multiple layers of features from tiny images.
- Kulkarni, V.; Kulkarni, M.; and Pant, A. 2020. Survey of personalization techniques for federated learning. In *2020 fourth world conference on smart trends in systems, security and sustainability (WorldS4)*, 794–797. IEEE.
- Lee, R.; Kim, M.; Li, D.; Qiu, X.; Hospedales, T.; Huszár, F.; and Lane, N. D. 2023. FedL2P: federated learning to personalize. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 14818–14836.
- Li, S.; Ye, F.; Fang, M.; Zhao, J.; Chan, Y.-H.; Ngai, E. C.-H.; and Voigt, T. 2024. Synergizing Foundation Models and Federated Learning: A Survey. *arXiv preprint arXiv:2406.12844*.
- Li, S.; Zhou, T.; Tian, X.; and Tao, D. 2022. Learning to collaborate in decentralized learning of personalized models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9766–9775.
- Li, T.; Hu, S.; Beirami, A.; and Smith, V. 2021a. Ditto: Fair and robust federated learning through personalization. In *International conference on machine learning*, 6357–6368. PMLR.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2: 429–450.
- Li, X.; JIANG, M.; Zhang, X.; Kamp, M.; and Dou, Q. 2021b. FedBN: Federated Learning on Non-IID Features via Local Batch Normalization. In *International Conference on Learning Representations*.
- Luo, J.; and Wu, S. 2022. Adapt to adaptation: Learning personalization for cross-silo federated learning. In *IJCAI: proceedings of the conference*, volume 2022, 2166. NIH Public Access.
- Ma, X.; Zhang, J.; Guo, S.; and Xu, W. 2022. Layer-wised model aggregation for personalized federated learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10092–10101.
- Mansilla, L.; Echeveste, R.; Milone, D. H.; and Ferrante, E. 2021. Domain generalization via gradient surgery. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6630–6638.
- Mansour, Y.; Mohri, M.; Ro, J.; and Suresh, A. T. 2020. Three approaches for personalization with applications to federated learning. *arXiv preprint arXiv:2002.10619*.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of

- deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Peng, D.; and Pan, S. J. 2023. Clustered task-aware meta-learning by learning from learning paths. *IEEE transactions on pattern analysis and machine intelligence*, 45(8): 9426–9438.
- Scott, J.; Zakerinia, H.; and Lampert, C. H. 2024. PeFLL: Personalized federated learning by learning to learn. In *The Twelfth International Conference on Learning Representations*.
- Shamsian, A.; Navon, A.; Fetaya, E.; and Chechik, G. 2021. Personalized federated learning using hypernetworks. In *International Conference on Machine Learning*, 9489–9502. PMLR.
- Smith, V.; Chiang, C.-K.; Sanjabi, M.; and Talwalkar, A. 2017. Federated multi-task learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 4427–4437.
- T Dinh, C.; Tran, N.; and Nguyen, J. 2020. Personalized federated learning with moreau envelopes. *Advances in neural information processing systems*, 33: 21394–21405.
- Woisetschläger, H.; Isenko, A.; Wang, S.; Mayer, R.; and Jacobsen, H.-A. 2024. A survey on efficient federated learning methods for foundation model training. *arXiv preprint arXiv:2401.04472*.
- Xie, C.; Huang, D.-A.; Chu, W.; Xu, D.; Xiao, C.; Li, B.; and Anandkumar, A. 2024. PerAda: Parameter-Efficient Federated Learning Personalization with Generalization Guarantees. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 23838–23848.
- Yang, F.-E.; Wang, C.-Y.; and Wang, Y.-C. F. 2023. Efficient model personalization in federated learning via client-specific prompt generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19159–19168.
- Yang, Y.; Long, G.; Shen, T.; Jiang, J.; and Blumenstein, M. 2024. Dual-Personalizing Adapter for Federated Foundation Models. *arXiv preprint arXiv:2403.19211*.
- Yi, L.; Yu, H.; Wang, G.; and Liu, X. 2023. FedLora: Model-heterogeneous personalized federated learning with lora tuning. *arXiv preprint arXiv:2310.13283*.
- Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020. Gradient surgery for multi-task learning. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 5824–5836.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*, 3987–3995. PMLR.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023. Fedala: Adaptive local aggregation for personalized federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 11237–11244.
- Zhang, M.; Sapra, K.; Fidler, S.; Yeung, S.; and Alvarez, J. M. 2021. Personalized Federated Learning with First Order Model Optimization. In *International Conference on Learning Representations*.
- Zhuang, W.; Chen, C.; and Lyu, L. 2023. When foundation model meets federated learning: Motivations, challenges, and future directions. *arXiv preprint arXiv:2306.15546*.