

Dual-Personalizing Adapter for Federated Foundation Models

Yiyuan Yang¹, Guodong Long¹, Tao Shen¹, Jing Jiang¹ and Michael Blumenstein¹

¹University of Technology Sydney

Abstract

Recently, foundation models, particularly large language models (LLMs), have demonstrated an impressive ability to adapt to various tasks by fine-tuning large amounts of instruction data. Notably, federated foundation models emerge as a privacy preservation method to fine-tune models collaboratively under federated learning (FL) settings by leveraging many distributed datasets with non-IID data. To alleviate communication and computation overhead, parameter-efficient methods are introduced for efficiency, and some research adapted personalization methods to federated foundation models for better user preferences alignment. However, a critical gap in existing research is the neglect of test-time distribution shifts in real-world applications. Therefore, to bridge this gap, we propose a new setting, termed test-time personalization, which not only concentrates on the targeted local task but also extends to other tasks that exhibit test-time distribution shifts. To address challenges in this new setting, we explore a simple yet effective solution to learn a comprehensive foundation model. Specifically, a dual-personalizing adapter architecture (FedDPA) is proposed, comprising a global adapter and a local adapter for addressing test-time distribution shifts and personalization, respectively. Additionally, we introduce an instance-wise dynamic weighting mechanism to optimize the balance between the global and local adapters, enhancing overall performance. The effectiveness of the proposed method has been evaluated on benchmark datasets across different NLP tasks.

1 Introduction

Recently, foundational models, especially the large language model (LLM) within the domain of natural language processing (NLP), have garnered significant interest [Brown *et al.*, 2020]. By utilizing vast amounts of data and sophisticated training algorithms, foundation models are endowed with a rich tapestry of generalized knowledge. To further refine these models for a more precise alignment with specific tasks

and user preferences, various fine-tuning methods have been explored. Notably, federated foundation models [Zhuang *et al.*, 2023; Yu *et al.*, 2023] represent an innovative approach, integrating federated learning frameworks for collaboratively fine-tuning the pre-trained foundation models by leveraging client-specific datasets to address privacy concerns.

Notwithstanding their considerable potential, direct fine-tuning of foundation models in FL incurs substantial computational and communication overhead due to the voluminous number of parameters encompassed. Parameter-efficient fine-tuning (PEFT) methods [Xu *et al.*, 2023a] present a promising alternative to mitigate these challenges. By selectively tuning and transmitting only a subset of parameters in FL, these methods seek to enhance efficiency. Among these PEFT methods, the adapter family [Hu *et al.*, 2023] (e.g., LoRA [Hu *et al.*, 2021]) stands out as one of the most popular methods in contemporary research of federated foundation models due to its flexibility compared with other methods.

Existing works have explored a spectrum of methodologies to address specific challenges when adapting the adapter-based PEFT methods within federated foundation models, including data heterogeneity [Babakniya *et al.*, 2023; Jiang *et al.*, 2023], communication overheads [Xu *et al.*, 2023b; Sun *et al.*, 2023], and so on. However, there are few studies delving into the personalization of federated foundation models for local clients, a crucial aspect in practical scenarios characterized by non-uniform data distribution across clients and the presence of diverse ability preferences among different clients. Furthermore, conventional personalized FL approaches [Tan *et al.*, 2022] predominantly concentrate on the specific targeted task for each client, often neglecting the challenges posed by test-time distribution shift when sometimes clients encounter other tasks during testing. For example, a client focusing on paper writing may sometimes need translation as an assistant. Given that a global model, trained on all mixed data, could not consistently excel in specifically targeted abilities as highlighted in [Wang *et al.*, 2023], there is a pressing need to devise methods for training personalized models that not only prioritize the performance on targeted abilities but also ensure satisfactory outcomes across other tasks, namely distribution shift on deploying federated foundation models.

To solve the aforementioned challenge, we introduce a brand-new setting close to real-world applications, noted as

test-time personalization, which follows: 1) each client needs to train a personalized model using its own data from a target task, and 2) during testing, each client’s personalized model needs to be capable to tackle the receiving new tasks (unseen in training) with different distributions (test-time distribution shift). This setting considers the test-time distribution shift scenario rather than only targeted tasks for personalization, which is essential in real applications of foundation models when clients may require assistance with tasks beyond their primary target task. Under this setting, a test-time personalized model should perform well on the personalized targeted task and perform comparably to other test-time tasks. To make this setting more challenging, we consider an extreme scenario where the proprietary tasks barely overlap across clients in training, but during the test, we consider test-time personalization in an ideal setting, where all tasks are included in all clients. As such, we take tasks of other clients with different distributions as test-time tasks for each client.

In test-time personalization, challenges become how to learn general knowledge from different clients for test-time tasks while prioritizing targeted ability for personalization. To address this issue, we explore a simple yet effective method, dubbed dual-personalizing adapter (FedDPA), where each client learns a global adapter to learn general knowledge from the aggregation for test-time tasks and maintains a local adapter for targeted ability personalization. During the inference phase, the local and global adapters are synergistically integrated to facilitate prediction, and an instance-wise dynamic weighting mechanism is proposed to autonomously adjudicate the proportional contribution of the local and global adapters. Experimental results demonstrate that our method achieves state-of-the-art performance on benchmarks. Our main contributions are summarized as follows:

- We propose a novel setting in personalized federated learning that emphasizes test-time distribution shifts in practical application scenarios, promoting comprehensive model learning during testing.
- We present a new method to realize test-time personalization, emphasizing the learning of both general and personalized knowledge for a more comprehensive model on various tasks.
- We conduct an exhaustive analysis using heterogeneous FL benchmarks across diverse NLP tasks. The empirical outcomes reveal that our method attains state-of-the-art performance, underscoring its superior test-time personalization capabilities in contrast to existing methods.

2 Related Work

2.1 Adapter-based PEFT Methods

Given the substantial computational and storage burdens associated with directly fine-tuning foundation models, the community has shifted towards embracing parameter-efficient methods [Xu *et al.*, 2023a], with the adapter family [Hu *et al.*, 2023] being a notable exemplar. According to different architectures, methods in the adapter family can be categorized into four types. The first one is prompt-based learning [Lester *et al.*, 2021; Li and Liang, 2021], which is aimed

at learning the continuous/soft prompt for discrete optimization. The second one is reparametrization-based methods [Hu *et al.*, 2021; Edalati *et al.*, 2022], achieving parameter efficiency by utilizing low-rank techniques to decompose the high-dimensional matrices. The third one is series Adapters [Houlsby *et al.*, 2019], which introduce additional learnable modules in a sequential manner within specific sublayers. The last one is parallel Adapters [He *et al.*, 2021], which focus on learning additional learnable modules in a parallel way with distinct sublayers. In this context, our exploration delves into the adapter-based PEFT methods of federated foundation models.

2.2 Federated Foundation Models

With the advent of foundation models, there has been a burgeoning interest [Zhuang *et al.*, 2023; Yu *et al.*, 2023] in integrating these models within the FL setting. Particularly, in light of the inherent computation and communication cost, recent work [Kuang *et al.*, 2023; Zhang *et al.*, 2023b] endeavors have delved deeper into integrating adapter-based parameter-efficient tuning (PEFT) methods with federated foundation models. Building upon this, a multitude of studies have emerged to navigate the challenges of incorporating federated foundation models with adapter-based PEFT methods. The paper [Zhang *et al.*, 2023a] stands at the forefront, initiating the integration of instruction tuning within federated LLM frameworks. Addressing data-related issues, the paper [Babakniya *et al.*, 2023] introduced a data-driven initialization approach to mitigate the primary challenges associated with LoRA in highly heterogeneous data scenarios. In addition, the research presented in [Jiang *et al.*, 2023] proposed a method to annotate unlabeled client-side data by harnessing the prowess of large models to address data scarcity concerns. To further optimize the communication and computational overheads associated with federated foundation models, the works [Xu *et al.*, 2023b; Sun *et al.*, 2023] emphasize advancing gradient-free optimization methods suitable for devices with limited memory and computing power. For personalization, paper [Yi *et al.*, 2023] focused on designing a specific training paradigm for LoRA to achieve more effective personalization in visual model-heterogeneous scenarios. Diverging from these approaches, our work delves into the realm of personalization with adapters in federated foundation models, extending the scope of research in this area.

2.3 Personalized Federated Learning

To address the necessity of personalization for individual clients, personalized Federated Learning (PFL) [Tan *et al.*, 2022], which aims at training to cater to individual client preferences and needs, is proposed. Broadly, existing PFL methods can be categorized into two primary types: fine-tuning the global model for personalization or learning additional personalized models. Research works [Fallah *et al.*, 2020; Collins *et al.*, 2021] in the first category fine-tuned the whole or part of the global model with each client’s local dataset for personalization. While research works [Li *et al.*, 2021a; Li *et al.*, 2021b] in the second category is to learn the additional personalized layers or model through local aggregation. Nonetheless, a prevalent limitation among these PFL ap-

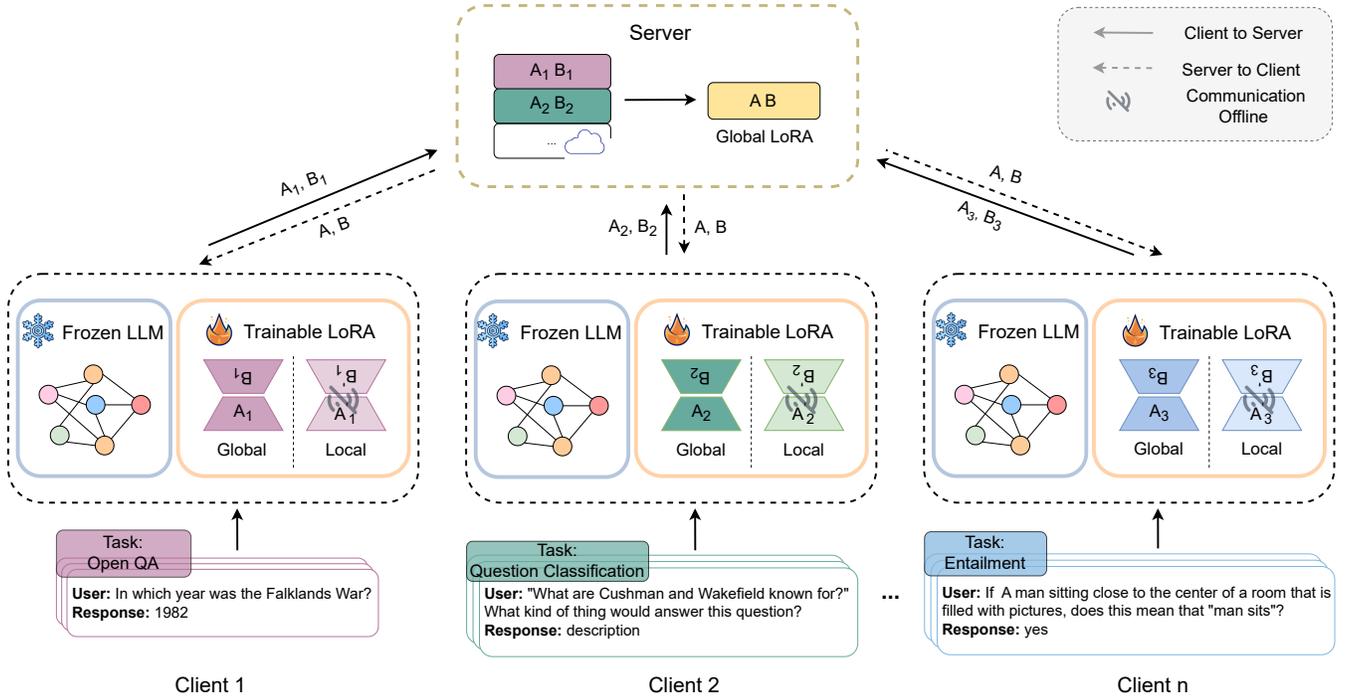


Figure 1: The overall framework of our proposed method. Each client contains a frozen LLM, a trainable global adapter (LoRA) and a trainable local adapter (LoRA) with a specific task, where the global adapter (LoRA) is for test-time tasks and the local adapter (LoRA) is for personalization. During the training, only the parameters of the global adapter (LoRA) are transmitted to the server for aggregation.

proaches is their concentrated focus on a specifically targeted task, often at the expense of performance when encountering test-time distribution shifts. Such a constrained approach might be suboptimal for practical applications that demand capability across other tasks. Therefore, we introduce the concept of test-time personalization to fill this gap.

3 Problem Definition

3.1 General FL Framework

In this paper, our method is built upon the general FL framework. In this framework, there are M clients, each possessing its distinct local privacy data \mathbb{D}^m , where m indexes a client. Concurrently, a server at the center level oversees the learning and maintenance of the aggregated model. For model aggregation across clients, we employ the FedAvg algorithm [McMahan *et al.*, 2017] without sacrificing generality, where the aggregated parameters θ are derived from a weighted sum of all clients' parameters θ^m based on the numbers of data. However, to address potential biases stemming from different numbers of tasks and data privacy concerns, our model aggregation strategy is based on the client number rather than the number of data. Consequently, the formulation can be represented as:

$$\theta = \sum_{i=1}^M \frac{1}{M} \theta^i \quad (1)$$

Local Tasks. In general FL framework, each local dataset, denoted as \mathbb{D}^m , is characterized by a set of data pairs repre-

sented as $\{(x_i^m, y_i^m) | i \in 1, \dots, N^m\}$, with N^m denoting the number of data. The objective for each local client is to minimize the empirical risk of all local data, which can be formulated as

$$\min_{\theta^m} \sum_{(x,y) \in \mathbb{D}^m} \mathcal{L}(x, y; \theta^m) \quad (2)$$

where \mathcal{L} is the loss function applied to each data instance and parameterized by θ^m .

3.2 Test-time Personalization

Here, we will introduce the training and testing phases separately for our test-time personalization setting.

In the training phase, the process and objective are the same as the general FL framework, where the model trains on the local data for empirical risk minimization. While in the testing phase, except for the test set \mathbb{D}_s driven from the same domain as training data, there are also some test sets \mathbb{D}_t under data distribution shifts $p_s(x, y) \neq p_t(x, y)$, and we call these datasets \mathbb{D}_t as test-time datasets. Therefore, the objective of the model should not only perform well on the test set \mathbb{D}_s but also have comparable results on the test-time dataset \mathbb{D}_t . This objective is consistent with the practical scenarios, since users primarily focus on the abilities they often utilize (abundant data available for training) and occasionally also introduce new tasks (limited to test data).

Remark. Let us clarify the concepts of ‘‘local dataset’’, ‘‘personalization’’ and ‘‘test-time’’. In our setting, the ‘‘local

dataset” refers to the data used in the training phase, serving as the source domain. “Personalization” denotes optimal performance in this source domain. The term “test-time” refers to the presence of additional data characterized by distribution shifts, used exclusively for testing. In our setting, test-time personalization aims to achieve the personalization while ensuring comparable performance on tasks encountered during the test-time phase.

4 Proposed Method

To simplify the illustration, we use LLM as the backbone of our proposed framework and adopt LoRA [Hu *et al.*, 2021] as the adapter-based PEFT method in our framework. The overall framework is easy to adapt to other types of backbone and other adapter-based PEFT methods. LoRA decomposes the training weight into a frozen weight θ , and a trainable weight derived through the multiplication of two low-rank weights $\Delta\theta = \Delta\theta^b\Delta\theta^a$. Given the inherent data heterogeneity intrinsic to federated learning environments, distinct NLP tasks are allocated to different clients. Consequently, within our FL framework, this data heterogeneity is manifest primarily in the distribution characteristics of the inputs across these diverse NLP tasks. Next, we will delve deeply into our proposed framework.

4.1 Overall Framework

In order to align with the application scenarios, we consider the test-time personalization setting, where the test-time distribution shifts data are tested. To simplify, we consider test-time personalization in an ideal setting, where all tasks are included in all clients. Therefore, for each client, the local task is taken as the primary task for personalization, while tasks from other clients are taken as the test-time tasks during the testing phase.

This framework raises two pivotal considerations: personalization and test-time distribution shifts. Since an aggregated global model at the server may not adequately address the specific tasks of each client, we introduce a local model to achieve personalization, with the premise that this personalization is anchored in maintaining local capabilities through training on each client’s local dataset. However, the aspect of distribution shifts during the testing phase remains a challenge, as conventional personalization methods often falter in these test-time tasks. Considering that data from other clients represents these test-time tasks, we can glean insights from the aggregated global model, which has already acquired a comprehensive understanding of all tasks through aggregation algorithms.

Therefore, we propose a dual-personalizing adapter (FedDPA) system for each client, based on the intuition that a global adapter targeting the test-time tasks and a local adapter for personalization are tuned together under a sophisticated federated learning algorithm. Additionally, we investigate two personalization methods to learn the local adapter, dubbed FedDPA-F and FedDPA-T respectively. As illustrated in Fig 1, the global adapter is obtained by conventional FL training and the local adapter is maintained locally by tuning on each local dataset. These two components are strategically combined to realize test-time personalization.

4.2 FL Training of Global Model

Addressing test-time distribution shifts necessitates a comprehensive acquisition of general knowledge. The conventional Federated Learning process is inherently designed to aggregate this general knowledge from a diverse range of tasks. Consequently, we utilize the adapter trained within the FL context as the global adapter for addressing test-time tasks, with more comprehensive details provided subsequently.

As illustrated in Fig 1, during the training stage, each client is endowed with a unique task and its corresponding local dataset \mathbb{D}^m , which is harnessed for model training. At each client, there consists of a frozen LLM model $f_m(x; \theta)$ with a global lightweight global adapter (LoRA) $\Delta\theta_m = \Delta\theta_m^b\Delta\theta_m^a$. This global adapter is used for aggregation by sending to the server. Notably, the server’s role is limited to computing the aggregated adapter $\Delta\theta$, thus obviating the need for maintaining a large-scale model. During each communication round k , the global adapter $\Delta\theta_m^k$ is seeded with the aggregated adapter $\Delta\theta^k$ from the central server, denoted as $\Delta\theta_m^k = \Delta\theta^k$. After this initialization, the model is trained on local datasets for a specified number of epochs, after which the updated adapter parameters are transmitted back to the server. Upon receipt of the adapter weights $\Delta\theta_m^k$ of all activated clients, the server employs FedAvg to aggregate these adapter weights. This process results in the formulation of the updated adapter weights $\Delta\theta^{k+1}$ for the next round. This iterative paradigm persists until training reaches its convergence. The overall objective can be formulated as:

$$\min_{\Delta\theta} \sum_{m=1}^M \sum_{(x,y) \in \mathbb{D}^m} \frac{1}{M} \mathcal{L}(x, y; \theta; \Delta\theta) \quad (3)$$

Remark. Other federated algorithms like FedProx [Li *et al.*, 2020] can also be applied with LoRA tuning of this global model learning. In this paper, we just take FedAVG as an example.

4.3 Personalization of Local Model

The primary purpose of test-time personalization lies in the concept of personalization. To actualize this, we integrate a local adapter into our approach. Drawing inspiration from prior research [Tan *et al.*, 2022], We introduce two methods for personalizing the local adapter: 1) directly fine-tuning the local adapter initialized by the global adapter, and 2) learning an additional local adapter during the training. In the first method, the local adapter is initialized by the global adapter, followed by subsequent fine-tuning using local datasets to facilitate personalization. The second method involves maintaining a separate local adapter locally during training without communication. At each stage of communication, this local adapter is initialized based on its previous state and then fine-tuned in conjunction with a frozen global adapter using local datasets, thereby facilitating the acquisition of personalized knowledge.

To be more specific, a local adapter (LoRA) $\Delta\hat{\theta} = \Delta\hat{\theta}^b\Delta\hat{\theta}^a$ is introduced. Thus, this model contains three components: a frozen LLM θ , a global adapter (LoRA) $\Delta\theta$ and a

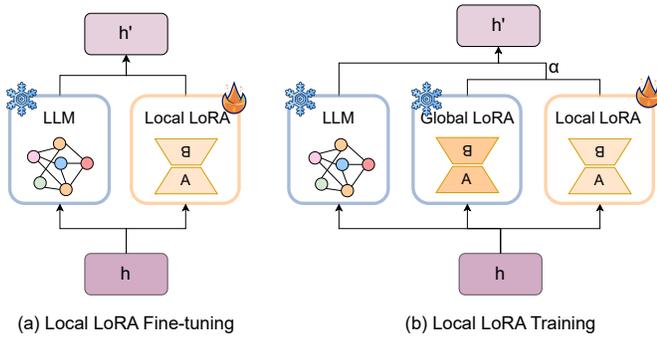


Figure 2: Frameworks of two personalized methods for the local adapter (LoRA).

local adapter (LoRA) $\Delta\hat{\theta}$. As delineated in Fig 2 (a), for the first method, after training, the local adapter is first initialized by the global adapter denoted as $\Delta\hat{\theta} = \Delta\theta$, then fine-tuned on local data to get the final local adapter. As shown in Fig 2 (b), for the second method, during each communication round in training for each adapter layer, upon receiving an input h , it simultaneously traverses the frozen LLM, the frozen global adapter and the local adapter. The process entails an initial fusion of the outputs from both the local and global adapters, followed by integration with the output of the LLM to yield the final result h' . A predefined weighting factor of α is employed to balance the contributions between the local and global adapters. The computational workflow in this configuration is mathematically formulated as follows:

$$h' = \theta h + ((1 - \alpha) \cdot \Delta\theta h + \alpha \cdot \Delta\hat{\theta} h) \quad (4)$$

4.4 Instance-wise Dynamic Weighting Mechanism

Upon training completion, we obtain a global adapter for test-time tasks and a local adapter for the personalized targeted task. The next question becomes how to combine them for subsequent inference tasks. This challenge is particularly pronounced in practical applications where the nature of the input instance—whether it pertains to the targeted task or test-time tasks—is not predetermined. Therefore, dynamically adjusting the weighting ratio between the global and local adapters for each instance becomes necessary. Considering the disparate data distributions that characterize test-time tasks and local tasks, coupled with the wealth of training instances specific to local tasks available to each client, we advocate for a novel approach to calculate the similarity between each incoming instance and the available training instances. This similarity metric serves as the key factor in determining the proportional contribution of the global and local adapters, thereby effectively guiding the combined weighting process and ensuring optimal performance across varying task scenarios.

Our proposed instance-wise dynamic weighting mechanism calculates the similarity between the input instance and local instances, using this metric to determine the appropriate weight balance for the global and local adapter combination. To facilitate this, the representation of each input instance is essential. Leveraging the robust capability of pre-trained

LLMs to abstract input sentences, we utilize the hidden states from the final layer of the LLM as the representation for similarity computation. Given that the LLM is decoder-based, with tokens attending only to preceding tokens, the embedding of the final token is considered representative of the entire input for similarity evaluation. Furthermore, to enhance the representation quality, the global adapter, which embodies general knowledge, is incorporated into this embedding process.

More specifically, during the inference stage, for each input instance x in a client, we randomly sample S instances $\{x_0, x_1, \dots, x_s\}$ from the local training dataset. These instances are then fed into the LLM, augmented with the global adapter, to obtain the last token’s embeddings from the final layer, denoted as w_x and $\{w_{x_0}, w_{x_1}, \dots, w_{x_s}\}$ respectively. Subsequently, we calculate the cosine similarity between the input representation w_x and each sampled local representation in $\{w_{x_0}, w_{x_1}, \dots, w_{x_s}\}$, resulting in a score range of $[0, 1]$. Finally, we average all scores to obtain the final result, represented as $\alpha = \sum_{i=0}^S \frac{1}{S} \cos(w_x, w_{x_i})$.

Through this methodology, the balancing of weights between the global and local adapters is dynamically adjusted for each test instance, ensuring the model not only tailors to the individual client’s specific needs but also benefits from the aggregated model’s collective knowledge across test-time tasks.

5 Experiment

5.1 Datasets and Baselines

In our experiment, we construct two federated datasets from Flan [Wei *et al.*, 2021]. Flan is a collection of diverse NLP tasks from over 60 datasets. For each dataset, different templates are used to transfer each example into an instruction for generative turning. In order to be better suitable for FL settings, we randomly select 8 NLP tasks from different datasets for each federated dataset and randomly select 300 examples for training and 200 examples for testing. ROGUE-1 is taken as a metric.

Here, we compare our methods with four baselines based on the same model architecture: centralized model, Local-finetuned model, FedIT [Zhang *et al.*, 2023a] and FedLoRA model [Yi *et al.*, 2023]. The centralized model is trained on all data of tasks in one center. The local-finetuned model infers that only local data are used to train the model without any communication with other clients or the server. Here, we adapt the training paradigm in FedLoRA [Yi *et al.*, 2023] to NLP tasks.

5.2 Implementation Details

We distribute the data between clients based on the NLP task for data heterogeneity. Since we select 8 NLP tasks, corresponding to $M = 8$ clients in our experiment. To better evaluate the effectiveness of our method, we assume that all clients are activated for every communication round and set the communication round K to 20. For local training, each client conducts 10 local epochs with a batch size of 32. During training, we adapt alpaca-LoRA as our base model and initialize it with LLaMA-7B. The rank of LoRA is set as

Methods	Federated Dataset 1								
	Para-phrase	Entail-ment	Structure to Text	Text For-matting	Linguistic Acc	Word Dis	Core-ference	Question CLS	Average
<i>Personalization</i>									
Centralized	77.00	82.00	72.58	96.59	70.50	63.50	77.59	89.00	78.60
FedIT	69.00	83.00	71.25	96.32	71.50	62.50	75.43	91.50	77.50
FedLoRA	77.50	84.00	71.49	96.69	73.50	65.00	75.27	92.00	79.43
Local-finetuned	74.50	80.00	73.71	97.36	75.00	54.50	68.55	89.50	76.64
FedDPA-F	79.00	84.50	72.06	96.90	72.00	65.00	73.86	92.50	79.48
FedDPA-T	80.50	84.50	72.79	96.51	73.50	62.00	77.93	94.00	80.22
<i>Test-Time Personalization</i>									
Local-finetuned	48.99	47.24	27.53	22.66	48.86	49.07	46.45	52.09	42.86
FedLoRA	75.56	76.55	75.21	74.94	76.16	74.64	74.99	76.97	75.63
FedDPA-F	78.10	77.36	77.18	76.98	77.11	76.23	76.84	77.19	77.12
FedDPA-T	76.20	75.51	76.19	75.63	74.86	74.60	74.77	75.96	75.47

Table 1: Personalization and test-time personalization results of different models on federated dataset 1. FedDPA-F represents the model with the local fine-tuning adapter and FedDPA-T represents the model with the local training adapter. Linguistic represents the linguistic acceptability task, Word Dis represents the word disambiguation task, and Question CLS represents the question classification task.

$r = 8$ and only applied to W_q and W_v . The updating weight of each client’s local LoRA during training is $\alpha = 0.5$ for federated dataset 1 and $\alpha = 0.3$ for federated dataset 2. We set $S = 5$ to select instances for the Instance-wise dynamic weighting mechanism.

5.3 Main Results

In this section, we represent the experiments of our methods compared with other baselines. This evaluation encompasses two primary facets: personalization (scores on targeted local tasks) and test-time personalization (average scores on all tasks including targeted local tasks and test-time tasks).

As evidenced in Table 1 and Table 2, our proposed dual-personalizing adapter methods (both fine-tuning and training) exhibit superior performance in personalization compared to other baseline models, which demonstrates the effectiveness of local adapter maintenance for enhancing performance on the targeted local task. For test-time personalization, the FedDPA-F method stands out as the most effective among all personalized models, which suggests that incorporating learning from the global adapter can be instrumental in adapting to test-time distribution shifts for a more comprehensive model achievement. In addition, it is noteworthy that while centralized or global models may yield higher average performances across all tasks, they fall short in excelling at specific tasks for personalization, aligning with the conclusions of the previous study [Wang *et al.*, 2023].

6 Analysis

6.1 Convergence Analysis

We present the convergence analysis of our methods in Figure 3. As illustrated in Figure 3 (a), we compared our methods with FedIT and FedLoRA for personalization, with the results showcasing the average performance on target local tasks across all clients. Notably, our methods exhibit a more rapid convergence compared to FedIT and achieve notable performance enhancements after five communication

rounds. Despite sharing similar trends with FedLoRA, our approaches, particularly the FedDPA-T, ultimately outperform in personalization. For a more granular insight into test-time personalization convergence, we contrast our approaches with other baselines, delineating average performance on all tasks, including each client’s targeted local and test-time tasks. Figure 3 (b) substantiates that our approaches demonstrate faster convergence rates, further bolstering the efficacy of our methods. Our methods demonstrate parallel trends with FedIT due to the benefit of general knowledge from the global adapter for test-time personalization.

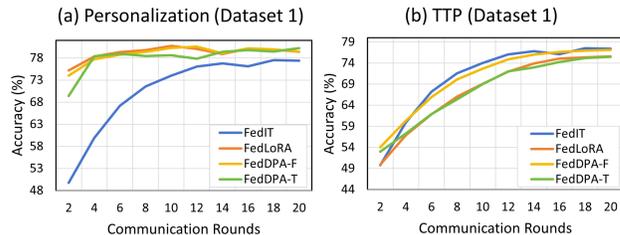


Figure 3: Average accuracy varies as communication rounds. TTP represents test-time personalization.

6.2 Ablation Study

Impact of Instance-Wise Dynamic Weighting Mechanism.

To explore the impact of the instance-wise dynamic weighting mechanism, we implemented experiments with FedDPA methods on different datasets. As shown in Table 3, the incorporation of an instance-wise dynamic weighting mechanism contributes significantly to enhancing performance in both personalization and test-time personalization scenarios. This enhancement is particularly pronounced for FedDPA-T.

Impact of Updating Weight α . In this study, we investigated the influence of the updating weight α during FedDPA-T training with its value $\alpha \in \{0.3, 0.5, 0.7\}$. As can be seen

Methods	Federated Dataset 2								
	Para -phrase	Common -sense	Entail -ment	Text For -matting	Summari -zation	Reading Com	Senti -ment	Open QA	Average
<i>Personalization</i>									
Centralized	87.00	64.67	77.00	90.65	29.12	76.00	72.50	76.17	71.64
FedIT	86.00	63.13	79.00	89.80	30.36	75.50	72.00	81.06	72.07
FedLoRA	87.00	64.12	84.50	89.52	27.13	76.50	73.50	79.62	72.74
Local-finetuned	75.00	53.51	81.00	91.28	27.51	69.00	72.50	79.31	68.64
FedDPA-F	88.00	64.80	84.25	89.82	29.58	78.50	72.00	80.89	73.48
FedDPA-T	90.50	70.54	82.00	91.81	30.75	81.00	75.00	91.07	75.33
<i>Test-Time Personalization</i>									
Local-finetuned	48.21	49.07	49.75	21.86	17.35	48.57	44.04	48.19	40.88
FedLoRA	69.60	71.64	71.09	71.28	65.63	68.89	70.32	70.44	69.86
FedDPA-F	71.64	72.28	72.42	72.39	71.12	70.46	71.00	71.82	71.64
FedDPA-T	71.63	72.66	71.20	72.58	70.58	69.21	70.67	71.62	71.27

Table 2: Personalization and test-time personalization results of different models on federated dataset 2. FedDPA-F represents the model with the local fine-tuning adapter and FedDPA-T represents the model with the local training adapter. Reading Com represents the reading comprehension task.

Methods	Auto	Fed Dataset 1		Fed Dataset 2	
		P	TTP	P	TTP
FedDPA-F	✗	79.06	76.97	73.17	71.70
	✓	79.48	77.12	73.48	71.64
FedDPA-T	✗	79.57	60.06	73.75	63.57
	✓	80.22	75.47	75.33	71.27

Table 3: Ablation study of instance-wise dynamic weighting mechanism (Auto). P represents personalization, and TTP represents test-time personalization.

Methods	α	Fed Dataset 1		Fed Dataset 2	
		P	TTP	P	TTP
FedDPA-T	0.3	79.69	75.85	75.33	71.27
	0.5	80.22	75.47	74.10	70.72
	0.7	79.88	75.01	74.04	69.95

Table 4: Ablation study of updating weight. P represents personalization, and TTP represents test-time personalization.

in Table 4, for test-time personalization, increasing updating weight α will decrease the performance due to the increased proportion of the local adapter in the model, while for personalization, different updating weights α are required for different datasets to achieve their optimal results.

Impact of Client Sample Rate. To explore how the number of participated clients impacts model performance, we implemented experiments with sample rate $\{0.2, 0.4, 0.6, 0.8, 1\}$. More specifically, we split the data of federated dataset 1 into 5 subsets for each task, where each subset has an equal number of training data and is assigned to one client. For each communication round, the server will select clients from each task based on the sample rate. As shown in Figure 4, as the client participant rates increase, model accuracy also increases as more participating clients provide more data for knowledge learning. Besides, FedDPA-

F performs better than FedDPA-T due to the possibility of overfitting when handling a small dataset,

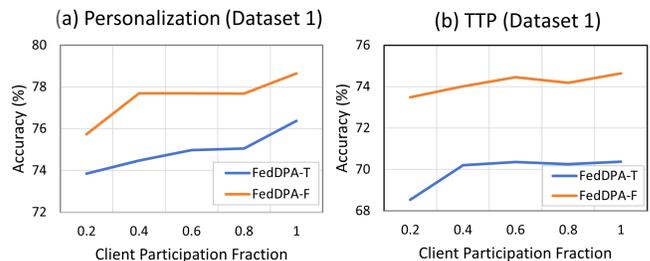


Figure 4: Average accuracy varies as different client participation fractions. TTP represents test-time personalization.

7 Conclusion

In this work, we propose a novel setting in federated foundation models by considering the test-time distribution shifts. In this setting, it is imperative for models to not only excel in performance on targeted local datasets but also to yield comparable results in test-time tasks. To address these challenges, we present a new method FedDPA, consisting of two core components: a global adapter to acquire general knowledge for test-time tasks, and a local adapter to learn personalized information for targeted local tasks. To further enhance the performance, an instance-wise dynamic weighting mechanism is introduced for balancing the local and global adapters during the inference. Compared with existing methods, our method achieves promising results on the constructed federated datasets with various NLP tasks, demonstrating its effectiveness.

References

- [Babakniya *et al.*, 2023] Sara Babakniya, Ahmed Roushdy Elkordy, Yahya H Ezzeldin, Qingfeng Liu, Kee-Bong Song, Mostafa El-Khomy, and Salman Avestimehr. Slora: Federated parameter efficient fine-tuning of language models. *arXiv preprint arXiv:2308.06522*, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [Collins *et al.*, 2021] 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.
- [Edalati *et al.*, 2022] Ali Edalati, Marzieh Tahaei, Ivan Kobzyev, Vahid Partovi Nia, James J Clark, and Mehdi Rezagholizadeh. Krona: Parameter efficient tuning with kronecker adapter. *arXiv preprint arXiv:2212.10650*, 2022.
- [Fallah *et al.*, 2020] 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.
- [He *et al.*, 2021] Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. Towards a unified view of parameter-efficient transfer learning. *arXiv preprint arXiv:2110.04366*, 2021.
- [Houlsby *et al.*, 2019] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pages 2790–2799. PMLR, 2019.
- [Hu *et al.*, 2021] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [Hu *et al.*, 2023] Zhiqiang Hu, Yihuai Lan, Lei Wang, Wanyu Xu, Ee-Peng Lim, Roy Ka-Wei Lee, Lidong Bing, and Soujanya Poria. Llm-adapters: An adapter family for parameter-efficient fine-tuning of large language models. *arXiv preprint arXiv:2304.01933*, 2023.
- [Jiang *et al.*, 2023] Jingang Jiang, Xiangyang Liu, and Chenyou Fan. Low-parameter federated learning with large language models. *arXiv preprint arXiv:2307.13896*, 2023.
- [Kuang *et al.*, 2023] Weirui Kuang, Bingchen Qian, Zitao Li, Daoyuan Chen, Dawei Gao, Xuchen Pan, Yuexiang Xie, Yaliang Li, Bolin Ding, and Jingren Zhou. Federatedscope-llm: A comprehensive package for fine-tuning large language models in federated learning. *arXiv preprint arXiv:2309.00363*, 2023.
- [Lester *et al.*, 2021] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- [Li and Liang, 2021] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- [Li *et al.*, 2020] 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.
- [Li *et al.*, 2021a] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pages 6357–6368. PMLR, 2021.
- [Li *et al.*, 2021b] Xiaoxiao Li, Meirui Jiang, Xiaofei Zhang, Michael Kamp, and Qi Dou. Fedbn: Federated learning on non-iid features via local batch normalization. *arXiv preprint arXiv:2102.07623*, 2021.
- [McMahan *et al.*, 2017] 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.
- [Sun *et al.*, 2023] Jingwei Sun, Ziyue Xu, Hongxu Yin, Dong Yang, Daguang Xu, Yiran Chen, and Holger R Roth. Fedbpt: Efficient federated black-box prompt tuning for large language models. *arXiv preprint arXiv:2310.01467*, 2023.
- [Tan *et al.*, 2022] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Wang *et al.*, 2023] Yizhong Wang, Hamish Ivison, Pradeep Dasigi, Jack Hessel, Tushar Khot, Khyathi Raghavi Chandu, David Wadden, Kelsey MacMillan, Noah A Smith, Iz Beltagy, et al. How far can camels go? exploring the state of instruction tuning on open resources. *arXiv preprint arXiv:2306.04751*, 2023.
- [Wei *et al.*, 2021] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- [Xu *et al.*, 2023a] Lingling Xu, Haoran Xie, Si-Zhao Joe Qin, Xiaohui Tao, and Fu Lee Wang. Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment. *arXiv preprint arXiv:2312.12148*, 2023.
- [Xu *et al.*, 2023b] Mengwei Xu, Yaorong Wu, Dongqi Cai, Xiang Li, and Shangguang Wang. Federated fine-tuning of billion-sized language models across mobile devices. *arXiv preprint arXiv:2308.13894*, 2023.

- [Yi *et al.*, 2023] Liping Yi, Han Yu, Gang Wang, and Xiaoguang Liu. Fedlora: Model-heterogeneous personalized federated learning with lora tuning. *arXiv preprint arXiv:2310.13283*, 2023.
- [Yu *et al.*, 2023] Sixing Yu, J Pablo Muñoz, and Ali Janesari. Federated foundation models: Privacy-preserving and collaborative learning for large models. *arXiv preprint arXiv:2305.11414*, 2023.
- [Zhang *et al.*, 2023a] Jianyi Zhang, Saeed Vahidian, Martin Kuo, Chunyuan Li, Ruiyi Zhang, Guoyin Wang, and Yiran Chen. Towards building the federated gpt: Federated instruction tuning. *arXiv preprint arXiv:2305.05644*, 2023.
- [Zhang *et al.*, 2023b] Zhuo Zhang, Yuanhang Yang, Yong Dai, Qifan Wang, Yue Yu, Lizhen Qu, and Zenglin Xu. Fedpetuning: When federated learning meets the parameter-efficient tuning methods of pre-trained language models. In *Annual Meeting of the Association of Computational Linguistics 2023*, pages 9963–9977. Association for Computational Linguistics (ACL), 2023.
- [Zhuang *et al.*, 2023] Weiming Zhuang, Chen Chen, and Lingjuan Lyu. When foundation model meets federated learning: Motivations, challenges, and future directions. *arXiv preprint arXiv:2306.15546*, 2023.

Appendix

A Implementation Details

A.1 Datasets

In this paper, we have developed two federated datasets derived from the Flan [Wei *et al.*, 2021], and details to construct these datasets are elucidated in this section. Flan encompasses a diverse array of NLP tasks, each comprising multiple datasets. To align with FL settings, we employed a stratified selection process, randomly choosing one dataset from each of the eight distinct tasks from Flan to form each federated dataset. In addition, to simulate client local data scarcity [McMahan *et al.*, 2017], we implemented a downsampling strategy, reducing the size of each selected dataset to 300 training instances and 200 testing instances. Consequently, each constructed federated dataset encompasses eight distinct NLP tasks, with a total of 300 training examples and 200 testing examples per task, culminating in a whole dataset comprising 2400 training examples and 1600 testing examples across all tasks. The specific tasks and datasets included in each federated dataset are cataloged in Table 5.

The NLP tasks within these datasets can be broadly divided into two types: generation tasks and classification tasks. To facilitate uniform processing by LLM, all tasks are converted into a generative format, employing distinct instructions for each dataset. Illustrative examples of these data for both classification and generation tasks are provided in Table 6. For the input of the LLM, we adopted a simple template, the details of which are delineated in Table 7.

Dataset Partitioning for Ablation Study. In our ablation study in section 6.2 examining the client sample rate to align with FL settings, we divided each task in our constructed federated datasets into five subsets, each comprising an equal number of training data. Based on our assumption that each client is associated with a single task, this division results in a total of 40 clients, with each client possessing a dataset of 60 training examples. To mimic real-world FL communication dynamics, we employed a randomized selection process for clients (subsets) within each task according to specified sample rates. Accordingly, for sample rates specified as $\{0.2, 0.4, 0.6, 0.8, 1\}$, we selected 1,2,3,4, and 5 clients (subsets) per task, leading to 8, 16, 24, 32, and 40 clients participating in federated communications, respectively. The evaluation phase involves computing the average results across these selected clients for each specified sample rate, which provides a comprehensive analysis of how client sample rates influence the performance of our method.

A.2 Baselines

In this section, detailed descriptions of the implementation for each baseline compared in this study will be provided:

- **Centralized model:** This model is formulated by aggregating all available data from various tasks at a single centralized center for training purposes, with 50 epochs to optimize.
- **Local-finetuned model:** This model trains independently of any external communication with other clients

or a central server. It is specifically trained on data pertaining to a single task, dedicating 50 epochs to optimize for task-specific performance without the influence of external data.

- **FedIT [Zhang *et al.*, 2023a]:** The FedIT model is the final aggregated global model derived from diverse local client datasets after training. It embodies the essence of collaborative learning inherent to federated learning, assimilating knowledge from a multitude of client-specific data sources.
- **FedLoRA model [Yi *et al.*, 2023]:** Here, we adapt the training paradigm in paper [Yi *et al.*, 2023] to NLP tasks by focusing on training the lightweight LoRA for aggregation while keeping the majority of the LLM parameters frozen. Subsequently, a personalized adaptation process is employed, where the globally aggregated LoRA undergoes further refinement on each local client’s dataset to tailor the learning outcomes to individual client needs.

B Additional Experiments

B.1 Instance-Wise Dynamic Weighting Mechanism Analysis

In this section, we further examine the impact of the instance-wise dynamic weighting mechanism, including the selected local instance number and the type of instance representation.

Impact of Instance Number S . In Section 4.4, the selection of S , representing the number of local instances for similarity calculation, is pivotal. To comprehensively evaluate the effect of varying the number of these instances, we conduct a series of experiments employing distinct local instance numbers, specifically $S \in \{1, 3, 5, 7, 9\}$. The accuracy results, as depicted in Figure 5, illustrate the dependency of model performance on different instance numbers S . As demonstrated in Figure 5 (a), in the context of personalization, it is observed that our models attain a plateau in accuracy when the instance number exceeds 5. This indicates a stabilization in model performance beyond this threshold of local instances. Furthermore, Figure 5 (b) delves into the realm of test-time personalization. The findings here reveal similar results, indicating that variations in the instance number do not markedly impact the model’s performance in test-time personalization.

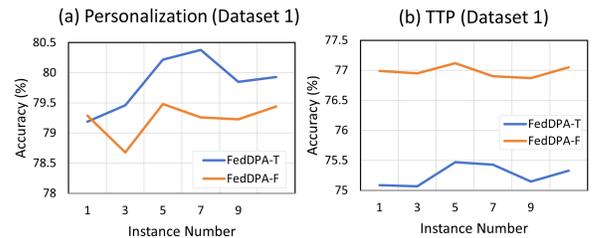


Figure 5: Average accuracy varies as different instance numbers. TTP represents test-time personalization.

Federated Dataset 1		Federated Dataset 2	
Task	Dataset	Task	Dataset
Paraphrase	glue_qqp	Paraphrase	paws_wiki
Entailment	snli	Commonsense	hellaswag
Structure to text	web_nlg_en	Entailment	qnli
Text formatting	fix_punct	Text formatting	word_segment
Linguistic acceptability	cola	Summarization	gigaword
Word disambiguation	wic	Reading comprehension	bool_q
Coreference	definite_pronoun_resolution	Sentiment	sentiment140
Question classification	trec	Open-domain QA	acr_easy

Table 5: Tasks and datasets of constructed federated dataset 1 and federated dataset 2.

Data Examples	
Input:	The father convinced his son that it is possible for him to one day become a knight, but he may never achieve such status coming from a peasant family. Who is "he"?
	OPTIONS: - The father - his son
Output:	His son
Input:	Police are seeking a former village chief in north china for allegedly killing his political rivals in an attack apparently motivated by local power plays, state press reported monday . Can you generate a short summary of the above paragraph?
Output:	Former chinese village head wanted for political murders

Table 6: Examples of data in our constructed federated datasets.

	Template
Prompt Input	Instruction: {instruction} Response:

Table 7: Prompt Template.

Methods	Emb	Fed Dataset 1	
		P	TTP
FedDPA-F	AVG	79.30	76.77
	LAST	79.48	77.12
FedDPA-T	AVG	79.65	73.36
	LAST	80.22	75.47

Table 8: Ablation study of instance representations (Emb). P represents personalization, and TTP represents test-time personalization. LAST represents using the embedding of the final token from the final hidden layer of LLM as instance representation, and AVG represents using the average embedding of all tokens from the final hidden layer of LLM as instance representation.

Impact of Instance Representation. In Section 4.4, our method entails utilizing the embedding of the final token from the last hidden layer of the LLM, denoted as 'LAST', as the input instance representation for the purpose of similarity calculation. In this exploration, we delve into another instance representation strategy, which involves employing the average embedding of all tokens from the final hidden layer of the LLM, herein referred to as 'AVG'. The comparative analysis, as presented in Table 8, demonstrates that employing the embedding of the last token yields superior performance relative to the strategy of averaging the embeddings of all tokens. This observed difference in performance can be attributed to the decoder structure inherent to LLMs, wherein the final token is capable of attending to all preceding tokens, thereby encapsulating comprehensive sentence-level information.

B.2 Model Scalability Analysis

In order to examine the effectiveness of model scalability, we conduct experiments based on a larger model, LLaMA-13B. The outcomes, as presented in Table 9, elucidate that larger models exhibit superior performance over their smaller counterparts across all personalization methods evaluated. Fur-

thermore, it is noteworthy that FedDPA-T surpasses FedDPA-F in terms of personalization and achieves comparable results in test-time personalization. This analysis underscores the inherent advantages of larger models in enhancing model performance, alongside the advance of the FedDPA-T approach in the context of personalization and adaptability to test-time conditions.

B.3 Communication and Computation Analysis

In this section, we undertake a detailed examination of both the communication and computation overhead associated with our proposed model in comparison to other baseline models. The results, as detailed in Table 10, delineate the communication and computation burdens imposed by various models. Given that these models are all based on the

Methods	Size	Fed Dataset 1	
		P	TTP
FedDPA-F	7B	79.48	77.12
	13B	81.52	80.55
FedDPA-T	7B	80.22	75.47
	13B	82.76	80.47

Table 9: Ablation study of model size. P represents personalization, and TTP represents test-time personalization.

Methods	Comm.Overhead	Comp.Overhead
FedIT	4.2M(0.06%)	0.277 TFLOPS
FedLoRA	4.2M(0.06%)	0.277 TFLOPS
FedDPA-F	4.2M(0.06%)	0.277 TFLOPS
FedDPA-T	4.2M(0.06%)	0.281 TFLOPS

Table 10: The communication and computation overhead of FedDPA and other baselines on Federated Dataset 1.

Methods	Time
FedLoRA	3.84s
FedDPA (w/o auto)	3.91s
FedDPA	4.13s

Table 11: Average inference time per instance. Auto represents the instance-wise dynamic weighting mechanism.

LoRA framework and exclusively transmit LoRA weights for aggregation (with our methods specifically transmitting only the global LoRA weights), they inherently sustain a minimal communication overhead. Regarding the computation overhead, the LoRA architecture permits the training of both local and global LoRAs in parallel, resulting in a marginal increase in computational demands for FedDPA-T. Conversely, FedDPA-F learns the local LoRA through an additional fine-tuning phase, thereby not imposing any additional computational overhead during the training phase.

Additionally, we have conducted an analysis of the inference time associated with our models. This examination involved calculating the average inference time per instance for FedLoRA, FedDPA without the instance-wise dynamic weighting mechanism, and FedDPA. As illustrated in Table 11, it is observed that our methods incur slightly higher inference time compared to FedLoRA. This marginal increase in inference time underscores the efficiency of our proposed methods, demonstrating that the enhanced performance and capabilities are achieved with a minimal impact on computational efficiency during inference.